DOCUMENT RESUME

ED 364 226 IR 054 714

TITLE	Advanced Network Technology. Background Paper.
INSTITUTION	Congress of the U.S., Washington, D.C. Office of Technology Assessment.
REPORT NO	ISBN-0-16-041805-4; OTA-BP-TCT-101
PUB DATE	Jun 93
NOTE	90p.; For an earlier background paper, see ED 336 066.
PUB TYPE	Reports - Evaluative/Feasibility (142)
EDRS PRICE	MF01/PC04 Plus Postage.
DESCRIPTORS	*Computer Networks; Futures (of Society); *Information Networks; *Information Technology; Models; Pilot Projects; *Research Utilization; Systems Development; *Technological Advancement
IDENTIFIERS	Information Infrastructure; Internet; *National Research and Education Network

ABSTRACT

This background paper analyzes technologies for tomorrow's information superhighways. Advanced networks will first be used to support scientists in their work, but will soon be deployed more widely in business, entertainment, health care, and education. Significant progress has been made toward the development of gigabit network technology since the basic characteristics of the design of broadband networks began to emerge in the mid-1980s. No insurmountable technological barriers to the gigabit National Research and Education Network (NREN) appear to exist, as work in the testbeds (i.e., initial testing programs) is demonstrating. Testbed networks model the configuration in which the technology is expected to be deployed, in that test sites are separated by realistic distances and realistic technological applications will be used. Testbed applications research helps researchers understand how the NREN can be used to achieve science goals and as a testbed in itself, demonstrating technology that can be deployed more widely. The following topics are reviewed: (1) the Internet; (2) broadband network technology; (3) gigabit research; and (4) application of testbed research. One table and 17 figures illustrate the discussion. Highlighted points are summarized in 10 boxes. (SLD)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED 364 226

BACKGROUND PAPER

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☐ This document has been reproduced as received from the person or organization originating it

☐ Minor changes have been made to improve reproduction quality

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy

dvanced Network Technology

Office of Technology Assessment

2

Technology Assessment Board of the 103d Congress

EDWARD M. KENNEDY,
DON SUNDQUIST,

SENATE

HOUSE

ROGER C. HERDMAN
Director

Technology Assessment Advisory Council

JAMES C. HUNT

MAY LENNON

BACKGROUND PAPER

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☐ This document has been reproduced as received from the person or organization originating it

☐ Minor changes have been made to improve reproduction quality

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy

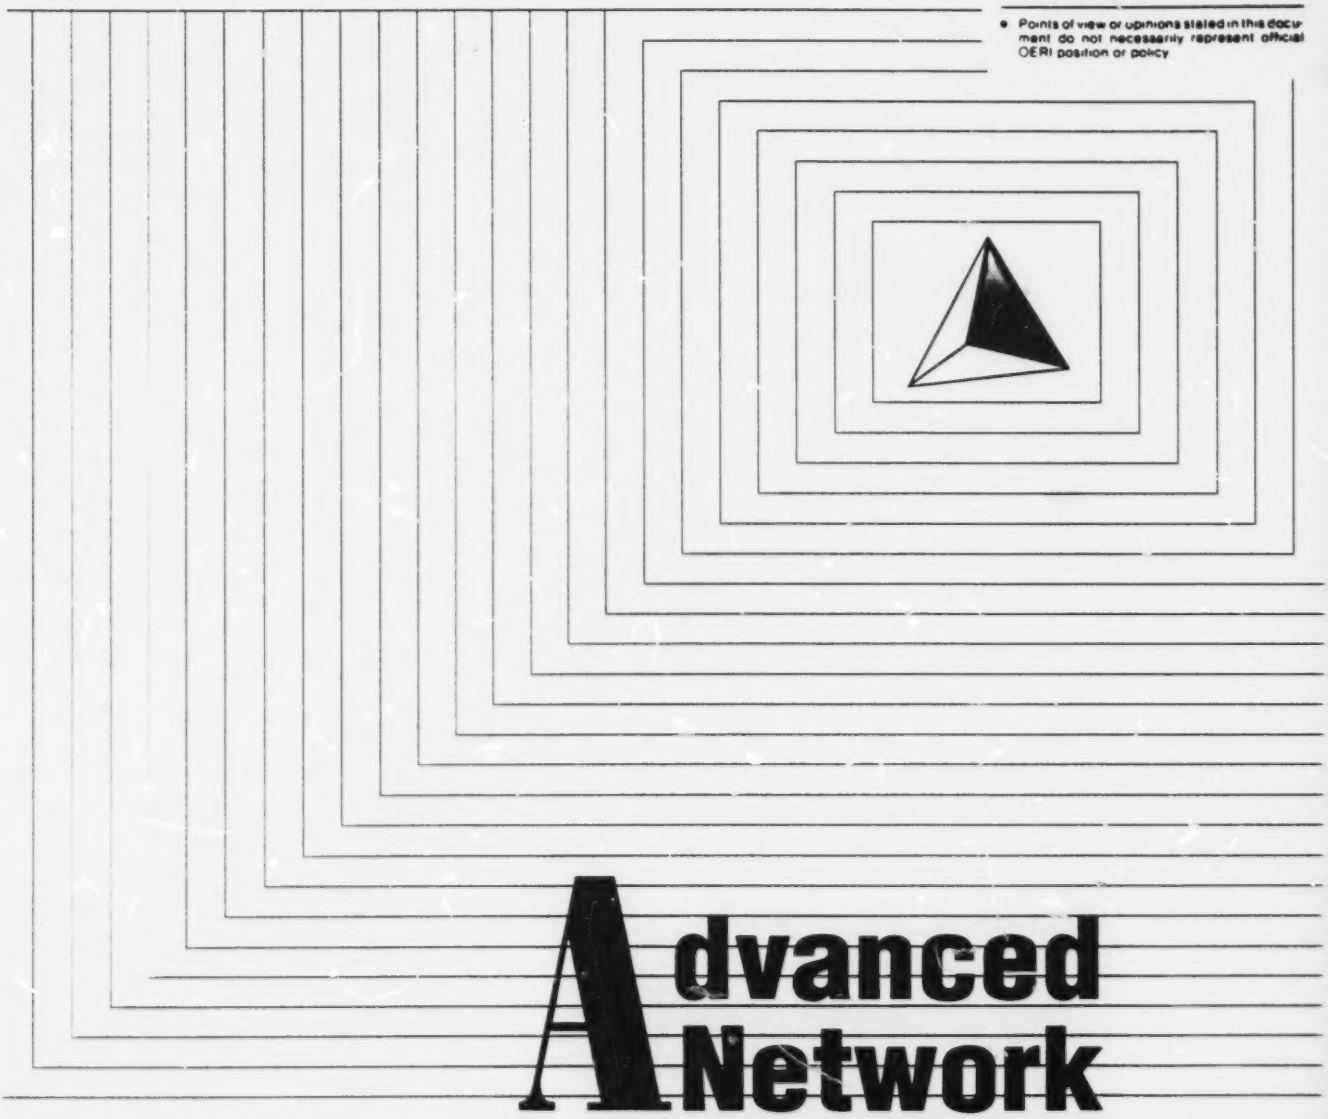

Advanced Network Technology

IR054714

BEST COPY AVAILABLE

Office of Technology Assessment

Recommended Citation:
U.S. Congress, Office of Technology Assessment, *Advanced Network Technology—Background Paper*, OTA-BP-TCT-101 (Washington, DC: U.S. Government Printing Office, June 1993).

For sale by the U.S. Government Printing Office
Superintendent of Documents, Mail Stop: SSOP, Washington, DC 20402-9328
ISBN 0-16-041805-4

Foreword

Computer networks are having dramatic impacts on our lives. What were once esoteric tools used only by scientists and engineers are becoming more widely used in schools, libraries, and businesses. At the same time, researchers are working to develop even more capable networks that promise to change fundamentally the way we communicate.

This background paper analyzes tecnnologies for tomorrow's information superhighways. Advanced networks will first be used to support scientists in their work, linking researchers to supercomputers, databases, and scientific instruments. As the new networks are deployed more widely, they will be used by a broader range of users for business, entertainment, health care, and education applications.

The background paper also describes six test networks that are being funded as part of the High Performance Computing and Communications Program. These test networks are a collaboration of government, industry, and academia, and allow researchers to try new approaches to network design and to attack a variety of research questions. Significant progress has been made in the development of technologies that will help achieve the goals of the High-Performance Computing Act of 1991.

This is the third publication from OTA's assessment on information technology and research, which was requested by the House Committee on Science, Space, and Technology and the Senate Committee on Commerce, Science, and Transportation. The first two background papers, *High Performance Computing & Networking for Science* and *Seeking Solutions: High-Performance Computing for Science*, were published in 1989 and 1991, respectively.

OTA appreciates the assistance of the National Science Foundation, the Advanced Research Projects Agency, the Department of Energy, the National Aeronautics and Space Administration, and many experts in industry and academia who reviewed or contributed to this document. The contents of this paper, however, are the sole responsibility of OTA.

Roger C. Herdman, Director

Reviewers

Rick Adams
CEO
UUNET Technologies

Robert Aiken
Department of Energy

Raymond Albers
Assistant Vice President
Technology Planning
Bell Atlantic

Alan Baratz
Applications Solutions Director
High Performance Computing and Communications
IBM

Adam Beguelin
Research Scientist
School of Computer Science
Carnegie Mellon University

Richard Binder
Principal Scientist
Corporation for National Research Initiatives

John Cavallini
Deputy Associate Director
Office of Scientific Computing
Department of Energy

Bruce Davie
Member of Technical Staff
Broadband Packet Switching Research
Bellcore

Darleen Fisher
Associate Program Manager
Division of Networking and Communications Research and Infrastructure
National Science Foundation

Linda Garcia
Senior Associate
Office of Technology Assessment

Tom Hausken
Analyst
Office of Technology Assessment

Milo Medin
Deputy Project Manager
NASA Science Internet Office
NASA

Paul Messina
Director
Caltech Concurrent Supercomputer Facility
California Institute of Technology

Craig Partridge
Senior Scientist
Bolt Beranek and Newman

Daniel Stevenson
Director
Communications Research
MCNC

Richard Thayer
Director
Federal Government Affairs
AT&T

Bo Thomas
Senior Federal Account Manager
Sprint

Philip Webre
Principal Analyst
Congressional Budget Office

Allan Weis
President
Advanced Network & Services

Joan Winston
Senior Analyst
Office of Technology Assessment

NOTE: OTA appreciates and is grateful for the valuable assistance and thoughtful critiques provided by the reviewers. The reviewers do not, however, necessarily approve, disapprove, or endorse this background paper OTA assumes full responsibility for the background paper and the accuracy of its contents.

Project Staff

ALAN BUZACOTT
Project Director

John Andelin
Assistant Director, OTA
Science, Information, and
Natural Resources Division

James W. Curlin
Program Manager, OTA
Telecommunication and
Computing
Technologies Program

Administrative Staff

Liz Emanuel, *Office Administrator*
Barbara Bradley, *Secretary*
Karolyn St. Clair, *PC Specialist*

Contents

1 Introduction and Summary, 1
Federal Support for Gigabit Networking, 1
Summary, 8

2 The Internet, 15
Applications, 21
Protocols, 24
Network Components, 26
The Internet and the Public Switched Network, 31

3 Broadband Network Technology, 35
Broadband Applications, 35
Fast Packet Networks, 42
Network Component Development-Current Status, 44
Application of Broadband Technologies, 47

4 Gigabit Research, 51
Research Objectives, 52
Testbed Progress, 56

5 Application of Testbed Research, 65
Application to the NREN, 65
Application to Other Networks, 70

INDEX, 75

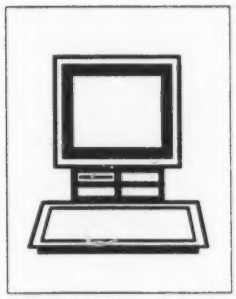

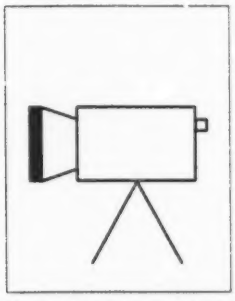

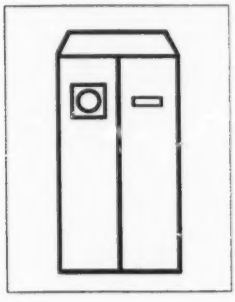

Introduction and Summary 1

The vision of the Nation's future telecommunications system is that of a broadband network (see box 1-A) that can support video, sound, data, and image communications. Toward this end, the High-Performance Computing Act of 1991 called for the Federal computer networks that connect universities and Federal laboratories to be upgraded to "gigabit networks" (see box 1-B) by 1996.[1] This background paper reviews technologies that may contribute to achieving this objective, and describes the six prototype gigabit networks or "testbeds" that are being funded as part of the Federal High Performance Computing and Communications Program. These prototype networks are intended to demonstrate new communications technologies, provide experience with the construction of advanced networks, and address some of the unresolved research questions.

FEDERAL SUPPORT FOR GIGABIT NETWORKING

The High Performance Computing and Communications Program (HPCC) is a multiagency program that supports research on advanced supercomputers, software, and networks.[2] In part, these technologies are being developed to attack the "Grand Challenges": science and engineering problems in climate change, chemistry, and other areas that can only be solved with powerful computer systems. Network research is one of four components of the HPCC program, and represents about 15 percent of the program's annual budget of close to $1 billion.[3]

The HPCC program funds the development of new communications technologies.

[1] High-Performance Computing Act of 1991 (HPCA), PL 102-194, Sec. 102(a).

[2] Office of Science and Technology Policy (OSTP), "Grand Challenges 1993: High Performance Computing and Communications," 1992.

[3] Ibid., p. 28.

Box 1-A—Broadband Networks

Computers and networks handle information as patterns of electronic or optical signals. Text, pictures, sound, video, and numerical data can then be stored on floppy disks, used in computations, and sent from computer to computer through a network. In *digital* computers or networks, the electronic or optical signals that represent information can take on one of two values, such as a high or a low voltage, which are usually thought of either as a "1" or a "0" (figure 1-A-1). These 1s and 0s are called *bits*.

Figure 1-A-1—Digital Data

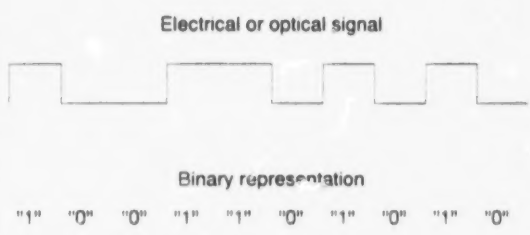

SOURCE: Office of Technology Assessment, 1993.

Different patterns of 1s and 0s are used to represent different kinds of data. In most computers, the letter "A" is represented by the pattern of electronic signals corresponding to "01000001." To represent images, different patterns of bits are used to represent different shades (from light to dark) and colors. Sound is represented in much the same way, except that the patterns of bits represent the intensity of sound at points in time.

The number of bits required to represent information depends on a number of factors. One factor is the quality of the representation. A good quality, high-resolution image would require more bits than a low-resolution image. Also, some kinds of information inherently require more bits in order to be represented accurately. A page of a book with only text might contain a few thousand characters, and could be represented with a few tens of thousands of bits. A page of image data, on the other hand, could require millions of bits.

Because images and video, which is a sequence of images, require many more bits to be represented accurately, they have strained the capabilities of computers and networks. Images take up too much space in a computer's memory, and take too long to be sent through a network to be practical. The new high-capacity network technologies described in this background paper have the ability to support two-way digital, image, and video communications in a more efficient manner.

Digital Networks

In the past, networks designed for video or sound used *analog* transmission. In the old analog telephone network, for example, the telephone's microphone converted the spoken sounds into an electrical signal whose

The other three components of the program target supercomputer design, software to solve the Grand Challenges, and research in computer science and mathematics.

The HPCC program is the most visible source of Federal funds for the development of new communications technology. The networking component of the program is divided into two parts: 1) research on gigabit network technology, and 2) developing a National Research and Education Network (NREN). The gigabit research program supports research on advanced network technology and the development of the six testbeds. The NREN program supports the deployment of an advanced network to improve and broaden network access for the research and education community. The High-Performance Computing Act of 1991 specifies that the NREN should operate at gigabit speeds by 1996, if technically possible.[4]

[4] HPCA, op. cit., footnote 1.

strength corresponded to the loudness of the sounds. This signal then traveled through the network's wires until it reached its destination, where it was used to make the telephone's speaker vibrate, recreating the spoken sounds.

Digital networks transmit information in digital form, as a series of bits. Digital networks are required for high-speed communications between computers—computers work with digital data. However, digital networks can also transmit real-world information such as sounds and pictures if special digital telephones or video cameras are used to represent the information in digital form. A digital telephone, for example, generates a series of patterns of 1s and 0s, corresponding to the loudness of the sounds. At the destination, these 1s and 0s are interpreted by the digital telephone and used to recreate the original sounds. Digital networks are quickly replacing analog networks. They are needed to transmit the growing amount of computer data. They also transmit voice and video information more cleanly, without interference and distortion. More importantly, digital networks allow a single network to carry all types of information. Today, separate networks are used for voice traffic (the telephone network), computer communications (data networks such as the internet) and video (broadcast or cable television or other specialized networks). Because these different kinds of information can all be represented in digital form, a single digital network can potentially be used to transmit all types of information. This is not the only requirement, however (see ch. 2 and ch. 3).

Broadband Networks

The capacity of a digital network is often described in terms of the number of bits that the network can transmit from place to place every second. A digital telephone network can transmit 64,000 bits every second. This is sufficient capacity to carry a telephone conversation with acceptable quality, but is not enough to carry video. Although some videotelephones can use regular telephone lines, users of videoconferencing systems usually prefer to use special services that can transmit at 384,000 bits per second or more. VCR-quality television needs about 1.5 million bits per second, and high-definition television needs about 20 million bits per second—about 300 times the capacity of a digital telephone line.

The capacity of a network, measured as the number of bits it can transmit every second, is called "bandwidth." Engineers often talk about "narrowband" networks, which are low bandwidth networks, and "broadband" networks, which are high bandwidth networks. The dividing line between the two is not always clear, and changes as technology evolves. Today, any kind of network that transmits at more than 100 million bits per second would definitely be considered a broadband network. Chapter 3 describes fiber optics and other technologies that will be used to build broadband networks.

SOURCE: Office of Technology Assessment, 1993.

Broadband networks such as the NREN will both improve the performance of existing applications and accommodate new types of applications. There will likely be a shift to image- and video-based communications, which are not adequately supported by currently deployed network technology. "Multimedia" applications that use images and video, as well as text and sound, look promising in a number of areas, e.g., education, health care, business, and entertainment. Broadband networks will also allow a closer coupling of the computers on a network; as the network is removed as a bottleneck, the computers will be able to form an integrated system that performs as a single, more powerful, computer.

Broadband networks will require a fundamental rethinking of network design. Several new concepts have been proposed and are being investigated by the testbeds. Fiber is a highly touted technology for constructing broadband networks, but it alone is not sufficient. Switches (see box 1-C) and the components that link computers to the network will have to be upgraded at the same time in order to keep pace with

Box 1-B—Gigabit Networks

Much of the research described in this background paper is aimed at the development of *gigabit networks*, broadband networks that can transmit data at one billion bits per second or more (a "gigabit" is one billion bits; "gigabit per second" is abbreviated as Gb/s or Gbps). This represents a 20-fold increase over the most capable links in the networks that currently serve the research and education community. The current National Science Foundation network uses links that transmit data at 45 million bits per second (megabits per second or Mb/s), and even this capacity has not been fully utilized because of bottlenecks in the network's switches. The development of a gigabit network is an ambitious target—most current industry technology planning targets broadband networks with lower bandwidths, in the 150 million bits per second range.

The basic outlines of the technology evolution of the DOD, NASA, DOE, and NSF networks that serve research and education were established in 1987 and 1989 reports issued by the Office of Science and Technology Policy. In the late 1980s, link bandwidths in the Federal networks were 1.5 Mb/s or less. The OSTP reports outlined a three-stage plan for the evolution of these networks to gigabit networks by the mid-to-late 1990s (see figure 1-B-1). The gigabit target was also specified by the High-Performance Computing Act of 1991. The OSTP report envisioned that each generation of technology would move from an experimental phase in the Federal networks to commercial service.

Figure 1-B-1—Timetable for the National Research and Education Network

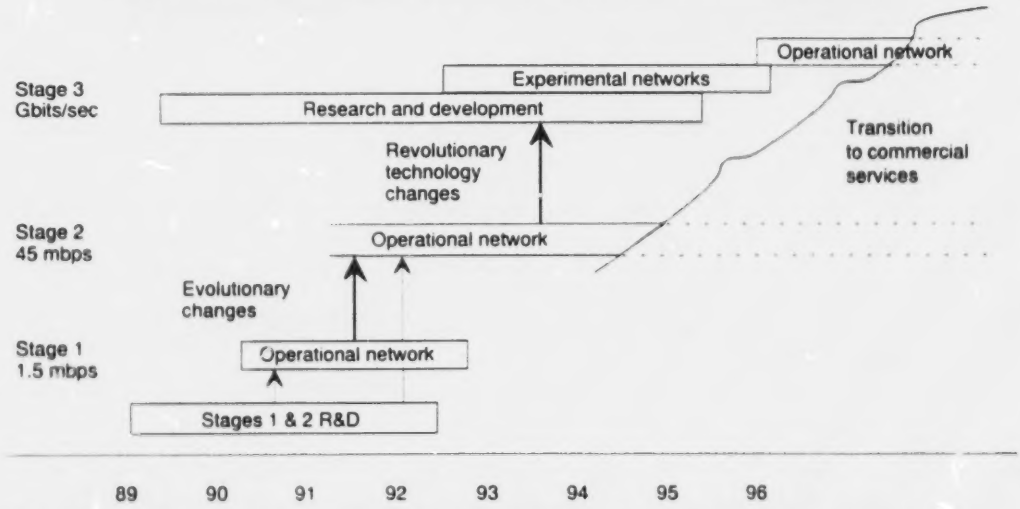

SOURCE: Office of Science and Technology Policy, "The Federal High Performance Computing Program," September 8, 1989.

Currently, the Federal agency networks are in the middle phases of the second stage, the operation of networks with 45 Mb/s links. At the same time, research and development for the third stage, the deployment of gigabit networks, is underway. In practice, the network capacity will not jump directly from 45 Mb/s to gigabit rates. The next step will be to 155 Mb/s, then to 622 Mb/s, and then to greater than one gigabit per second. The bandwidths used in computer networks (1.5 Mb/s, 45 Mb/s, 155 Mb/s, and 622 Mb/s) correspond to standards chosen by manufacturers of transmission equipment.

SOURCES: Office of Science and Technology Policy (OSTP), "A Research and Development Strategy for High Performance Computing," Nov. 20, 1987; OSTP, "The Federal High-Performance Computing Program," Sept. 8, 1989; High-Performance Computing Act of 1991 (HPCA), Public Law 102-194, Sec. 102(a).

Box 1-C—Computer Network Components

A computer network has three main components: computers, links, and switches (figure 1-C-1). The web of links and switches carry data between the computers. Links are made of copper (either "twisted pair" or "coaxial cable") or fiber optics. Transmission equipment at each end of the fiber or copper generates the electrical or optical signals. There are also satellite and microwave links that send radio waves through the air. Fiber has several advantages over other types of links—most notably its very high bandwidth. The fiber optic links needed for gigabit networks are already commercially available. However, gigabit networks will not be deployed until research issues in other network components are addressed.

Figure 1-C-1—A Simple Computer Network

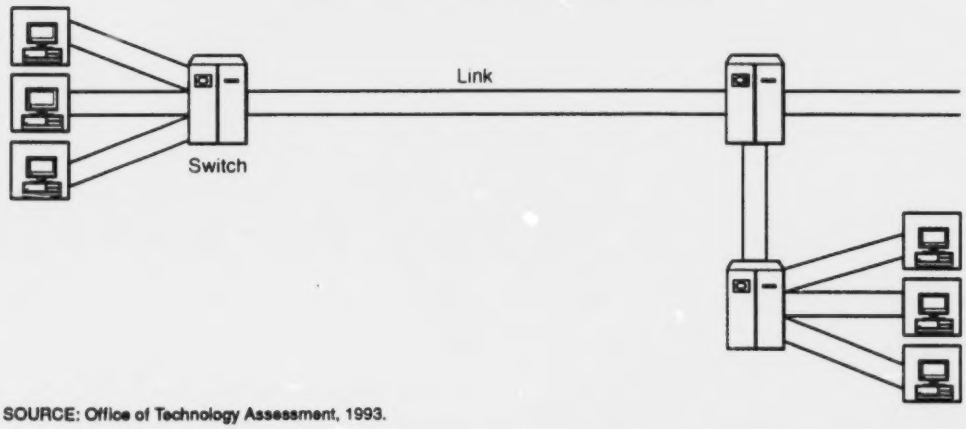

SOURCE: Office of Technology Assessment, 1993.

For example, new high-capacity switches are needed to keep pace with the higher bandwidth of fiber optic links. Just as railroad switches direct trains from track to track, the switches in computer networks direct information from link to link. As the information travels through the network, the switches decide which link it will have to traverse next in order to reach its destination. The rules by which the switches and users' computers coordinate the transmission of information through the network are called *protocols*.

While most computer networks are limited in their ability to carry high-bandwidth signals such as video, cable television networks are widely used to distribute television signals to homes. However, cable networks usually do not have switches. For this reason, they only permit one-way communications: the signal is simply broadcast to everyone on the network. Much of the network research today is devoted to the development of switches that would allow networks to support two-way, high-bandwidth communications.

SOURCE: Office of Technology Assessment, 1993.

the faster flow of data. Broadband networks will be more than simply higher bandwidth versions of today's networks, however. Networks will also be redesigned so that a single type of network can carry video, sound, data, and image services. The existing telephone and data networks do not have sufficient flexibility to carry all types of information efficiently.

The NREN

One objective for the NREN is that it serve as an enabling technology for science and engineer-

ing research.[5] The gigabit NREN will be able to handle the very large data sets generated by supercomputers. Scientists could use the gigabit NREN to support "visualization," the use of a computer-generated picture to represent data in image form. For example, ocean temperatures computed by a climate model could be represented by different colors superimposed on a map of the world, instead of a list of numbers. Visualization is an essential technique for understanding the results of a simulation. Currently, much of the data collected in experiments and computed by simulations goes unused because of the time needed to compute images on conventional computers. Supercomputers could perform the computations more quickly, but few laboratories have supercomputers. With a high-speed network, a scientist could send the data to a distant supercomputer, which would be able to quickly compute the images and send them back through the network for display on the scientist's computer.

Broadband networks will require a fundamental rethinking of network design.

A second objective for the NREN program is that it demonstrate and test advanced broadband communications technologies before they are deployed in commercial networks. The NREN program will upgrade federally supported networks such as the National Science Foundation's NSFNET, the Department of Energy's Energy Sciences Network (ESnet), and the National Aeronautics and Space Administration's NASA Science Internet (NSI).[6] These networks form the core of the "Internet," a larger collection of interconnected networks that provides electronic mail services and access to databases and supercomputers for users in all parts of the United States and around the world.[7] During 1992, Federal agencies announced plans for upgrading their current networks as part of the NREN program.[8]

The NREN program can be viewed as a continuation and expansion of the Federal support that created the Internet. The Internet's technology evolved from that of the Arpanet, a research project of the Advanced Research Projects Agency. Beginning in 1969, the Arpanet served to demonstrate the then-new technology of "packet switching." Packet switched networks were able to support computer communications applications that could not be efficiently accommodated by the telephone network's "circuit switched" technology (see ch. 2, p. 29). Packet switched networks are now widely deployed, Internet services are being offered by the private sector, and the Internet protocols are becoming world standards. In much the same way, the NREN program is intended to catalyze the deployment of a new generation of network technology.

Past government programs have also been successful in broadening access to networks for the larger research and education community. The Internet is increasingly essential to users in the academic community beyond the original core group of users in engineering and computer science. It is now estimated that over 600 colleges

[5] For a description of the goals and characteristics of the NREN see HPCA, op. cit., footnote 1, Sec. 102(a)-(c); OSTP, op. cit., footnote 2, p. 18; U.S. Congress, Office of Technology Assessment, *High Performance Computing & Networking for Science*, OTA-BP-CIT-59 (Washington, DC: U.S. Government Printing Office, September 1989), p. 25.

[6] OSTP, op. cit., footnote 2, p. 18; Office of Science and Technology Policy, "The National Research and Education Network Program: A Report to Congress," December 1992, p. 2.

[7] Robert E. Calem, "The Network of All Networks," *The New York Times*, Dec. 6, 1992, p. F12.

[8] National Science Foundation, "Public Draft: Network Access Point Manager/Routing Authority and Very High Speed Backbone Network Services Provider for NSFNET and the NREN Program," June 12, 1992; James F. Leighton, Manger of Networking and Engineering, National Energy Research Supercomputer Center, Lawrence Livermore National Laboratory, "ESnet Fast-Packet Services Requirements Specification Document," Feb. 20, 1992.

and universities and an estimated 1,000 high schools are connected to the Internet.[9] As the Internet's user community becomes more diverse, there is a growing need for simplifying the applications and their user interfaces.

This background paper primarily describes gigabit NREN applications and network technologies. There are, however, several controversial policy issues related to the NREN program.[10] First, the scope of the NREN is uncertain. As a key component of the HPCC program, a clear role of the NREN is to serve scientists and engineers at Federal laboratories, supercomputer centers, and major research universities. This objective will be met primarily by upgrading the networks operated by the National Science Foundation (NSF), Department of Energy (DOE), and the National Aeronautics and Space Administration (NASA). However, there are several different visions of the extent to which the NREN program should also serve a broader academic community, such as libraries and schools.

A second major issue concerns the "commercialization" of the NREN. The NREN will develop from the current Internet, which is increasingly used by government and businesses, not only by the research and education community. Several new commercial providers have emerged to offer Internet services to this market, which is not served by Federal agency networks. One of the goals of the NREN program is to continue this commercialization process, while at the same time achieving the science and network research goals of the NREN program. There has been considerable uncertainty about the mechanisms by which this objective is to be achieved.

The High-Performance Computing Act does not clearly specify the scope of the NREN or the mechanism for commercialization. NSF has had to address these issues in the course of developing a plan for the development of its network, which will be a central component of the NREN. These debates have slowed considerably the process by which NSF will select the companies that will operate its network. NSF's original plan, released in the summer of 1992, is undergoing significant revisions (see box 5-A). As of May, 1993, a new plan had not been issued. It is increasingly unlikely that NSF will be able to deploy its next-generation network by the Spring of 1994, as was originally planned.

In addition, the growing commercial importance of networking is leading to greater scrutiny of the agencies' choices of contractors to operate their NREN networks. DOE selected a contractor for its component of the NREN in the summer of 1992, planning to deploy the new network in mid-1993. However, a losing bidder protested DOE's selection to the General Accounting Office (GAO). In March, 1993, GAO overturned DOE's choice of contractor and recommended that DOE revise its solicitation, conduct discussions with potential contractors, and allow contractors a new opportunity to bid.[11] DOE has

[9] Darleen Fisher, Associate Program Manager, National Science Foundation, personal communication, Feb. 11, 1993.

[10] For issues related to the NREN program, see Hearings before the House Subcommittee on Science, Mar. 12, 1992, Serial No. 120.

[11] The dispute concerned the parties' interpretation of certain provisions in DOE's Request for Proposals (RFP). AT&T protested DOE's selection of Sprint to be the contractor for the DOE network, arguing sucessfully that the RFP had specified more fully-developed switches than had been proposed by Sprint as part of its bid. GAO ruled that the switches that Sprint planned to use did not comply with a provision in the RFP that proposals had to "conclusively demonstrate current availability of the required end-to-end operational capability." DOE, by contrast, was satisfied that the switches had been developed to the level envisioned by the RFP and were appropriate to a program designed to explore leading-edge technology.

DOE's RFP had specified the use of "cell relay" technology, which is the basis for both Asynchronous Transfer Mode (ATM) and Switched Multimegabit Data Service (SMDS) services. ATM is expected to play an important role in the future development of computer networking and the telecommunications industry, while SMDS is viewed primarily as an intermediate step towards ATM. DOE selected Sprint in large part because Sprint proposed to begin ATM services immediately, while AT&T bid a service based on SMDS and evolving to ATM only in 1994. Early deployment of ATM would have provided a valuable opportunity to evaluate and demonstrate a key telecommunications industry technology. Comptroller General of the United States, Decision in the Matter of AT&T, File B-250516.3, March 30, 1993.

asked GAO to reconsider its decision. The DOE example raises questions about the effect of government procurement procedures on the ability of federal agencies to act as pioneers of leading-edge network technology. The additional time that would be required to comply with GAO's recommendations, added to the seven-month GAO process, would delay deployment of DOE's network by over a year.

The Testbeds

The HPCC program's six gigabit testbeds (table 1-1) are intended to demonstrate emerging high-speed network technologies and address unresolved research questions. While each testbed involves a different research team and is emphasizing different topics, there is similarity in their approach. The testbeds typically consist of a high-speed network connecting three or four sites —universities, industry laboratories, supercomputer centers, and Federal laboratories—with high-bandwidth optical fiber. Located at each of the testbed sites are computers, prototype switches, and other network components. Each research group has both network and applications researchers—the applications will be used to test different approaches to network design.

Significant progress has been made toward the development of gigabit technology.

The testbed program is administered by NSF and the Advanced Research Projects Agency[12] (ARPA). Five of the testbeds are jointly funded for 3 years by NSF and ARPA under a cooperative agreement with the Corporation for National Research Initiatives (CNRI). The principals of CNRI, a nonprofit organization, played significant roles in the development of both the Arpanet and the Internet.[13] CNRI is responsible for organizing the testbeds and coordinating their progress. Funding for the testbeds is modest, when compared to their visibility and the overall HPCC budget. The cooperative agreement with CNRI is for $15.8 million over 3 years. Most of the cost of building the networks has been borne by industry, in the form of contributions of transmission capacity, prototype switches, and research personnel.

The testbeds are investigating the use of advanced network technology to match the needs of the NREN. There is an emphasis on delivering the highest bandwidths possible to the users and demonstrating the range of applications that would be used by leading-edge users of the NREN. Most of these applications are supercomputer-based. For example, some applications use the network to link several supercomputers, allowing their combined processing power to compute complex simulations more rapidly. Many of the applications being investigated also use the network to enable visualization of the results of simulations or experiments.

Initially, only a few users would have computers powerful enough to need a gigabit network. However, the processing power of lower cost workstations and ordinary desktop computers is likely to continue to increase rapidly, as a result of advances in microprocessor technology. Gigabit networks and the lessons learned from the testbeds will then be used more widely.

SUMMARY

Progress

Significant progress has been made toward the development of gigabit network technology since 1987, when the Office of Science and Technology Policy (OSTP) noted that considerable research would be needed to determine the design of

[12] Formerly the Defense Advanced Research Projects Agency (DARPA).

[13] Dr. Robert E. Kahn is President of CNRI; Dr. Vinton G. Cerf is Vice President.

330-073 O - 93 - 2 QL 3

Table 1-1—Gigabit Testbed Participants

Testbed	Location	Industry	Federal laboratories	Supercomputer centers	Universities & other
AURORA	Northeast	IBM Bellcore Bell Atlantic NYNEX MCI			MIT University of Pennsylvania
BLANCA	Nationwide	AT&T	Lawrence Berkeley Laboratory	National Center for Supercomputing Applicatons	University of Illinois University of Wisconsin University of California-Berkeley
CASA	Southwest	MCI Pacific Bell U.S. West	Jet Propulsion Laboratory Los Alamos National Laboratory	San Diego Supercomputer Center	California Institute of Technology
NECTAR	Pittsburgh	Bellcore Bell Atlantic		Pittsburgh Supercomputer Center	Carnegie Mellon University
VISTAnet	North Carolina	BellSouth GTE		North Carolina Supercomputer Center (at MCNC)	University of North Carolina-Chapel Hill North Carolina State University MCNC
MAGIC	South Dakota Kansas Minnesota	Sprint MITRE Digital Equipment Corp. Southwestern Bell Northern Telecom Split Rock Telecom SRI International	U.S. Army Future Battle Laboratory U.S. Army High-Performance Computing Research Center U.S. Geological Survey Lawrence Berkeley Laboratory	Minnesota Supercomputer Center	University of Kansas

SOURCE: Corporation for National Research Initiatives (CNRI), Advanced Research Projects Agency (ARPA).

gigabit networks.[14] There has been growing consensus within the technical community on many issues, and the development of the optical fiber links, switches, and other network components is underway. The testbeds represent the next step in the research—integrating the hardware and software components into a working network system and testing it with applications.

The basic characteristics of the design of broadband networks began to emerge in the mid-1980s, supported by the results of simulations and small-scale experiments. Researchers' objective was to develop networks that could support high bandwidths and were also sufficiently flexible to support a range of services. One characteristic of these networks is the use of optical fiber links, which have the necessary capacity to support many new services, including bandwidth-intensive video- and image-based applications. The second major characteristic of the proposed designs for advanced networks is the use of "fast packet switches," a new type of switch that has both the processing power to keep up with increases in link bandwidth and the flexibility to support several kinds of services.

The testbeds have established a useful model for network research.

As these ideas began to emerge, computer and telecommunications companies initiated the development of the network components required for broadband networks. There appear to be no significant technological barriers to the development of the components required for the gigabit NREN. Transmission equipment of the type that would be required for the gigabit NREN is already becoming available commercially and is being used in the testbeds. Some fast packet switches are also becoming commercially available. Versions of these switches that operate at gigabit rates are in prototype form and will be incorporated in the testbeds over the coming year.

The testbeds are looking to the next step in the research—the development of test networks. This is a systems integration task—developing the individual components is only part of the process of building an advanced network. There is often much to be learned about making the components work together and solving unforeseen problems in the implementation. In addition, there are research questions that can only be investigated with a realistic test network. The testbeds will provide a way to test various proposed approaches to network design.

Progress on the testbeds has been slower than expected, due to delays in making the transmission equipment available and in completing work on the switches and other components. Switches are complex systems, requiring the fabrication of numerous electronic circuits. It was originally hoped that the optical fiber links could be deployed and the gigabit switches and other components finished in time to have a year to experiment with the working testbed networks before the end of the program in mid-1993. It now appears that the testbeds will not be operational until the third quarter of 1993. The testbed program has been extended to permit a year's research on the testbed facilities once they become operational.

Testbed Concept

The testbeds have established a useful model for network research. The design and construction of a test network fills a gap between the earlier stages of the network research—small scale experiments and component development—and the deployment of the technology in production

[14] Office of Science and Technology Policy, "A Research and Development Strategy for High Performance Computing," Nov. 20, 1987, p. 21.

networks. The testbed networks model the configuration in which the technology is expected to be deployed—the test sites are separated by realistic distances and the networks will be tested with applications of the type expected to be used in the gigabit NREN. In addition, the participants in the testbeds will play important roles when the networks are deployed.

The testbed research contributes in a number of ways to a knowledge base that reduces the risks involved in deploying advanced network technology. First, there are a number of research issues that are difficult to address without a working network that can be used to try different approaches. Second, the systems integration process provides experience that can be applied when the production network is constructed. In many ways the experience gained in the process of getting the testbeds to work will be as valuable as any research done with the operational testbeds. Third, the testbeds serve to demonstrate the utility of the technology, which serves to create interest among potential users and commercial network providers.

The relatively small amount of government money invested has been used primarily to organize and manage the testbeds and to encourage academic involvement. The testbeds have mainly drawn on other government and industry investment. The organization of the testbeds as a collaborative effort of government, academic, and industry groups is essential, because of the many disciplines required to build and test a network. Industry has contributed expertise in a number of areas. For example, it would be too difficult and expensive for academic researchers to develop the high-speed electronics needed for the switches and other components. Academic researchers are involved in the Internet community, and have contributed ideas for new protocols and applications. Other applications work has come from a number of scientific disciplines and the supercomputer community.

One of CNRI's main contributions was to encourage the involvement of the telecommunications carriers in the testbeds. The transmission facilities required for the testbeds are expensive because of the long distances between the testbed sites and the demands for very high bandwidth. Most experimental work in the past was on small scale networks in a laboratory, due to the prohibitive cost of linking distant test sites. However, the carriers are installing the required transmission capacity and making it available to the testbeds at no cost. All three major interexchange carriers (AT&T, MCI, and Sprint), and most of the Regional Bell Operating Companies (RBOCs) are playing a role in the testbeds.

The testbed research overlaps with industry priorities in some areas and not in others. The basic design of the networks—the types of switches and transmission equipment—reflects emerging industry concepts. However, much of the research agenda focuses on higher bandwidths and more specialized applications than will be used with commercial broadband networks in the near term. Only a few users will use the types of supercomputer-based applications being emphasized by the testbeds. Of greater near-term commercial importance to industry are medium bandwidth "multimedia" applications that require more bandwidth than can be supported by current networks, but significantly less than the gigabit speeds required by the supercomputer community.

Application of Testbed Research

The testbed research is applicable both to the NREN and to other networks. The NREN will serve only the research and education community and is best viewed as only part of the broader national information infrastructure.[15] The scope of the national information infrastructure will in-

[15] For one view of the relationship between the NREN and the "National Information Infrastructure," see Michael M. Roberts, "Positioning the National Research and Education Network," *EDUCOM Review*, vol. 26, No. 3, summer 1991, pp. 11-13.

clude both the United States' part of the Internet and a wide array of other services offered by the computer and information industries, the carriers, the cable television industry, and others.

APPLICATION TO NREN

During 1992, DOE, NASA, and NSF published plans for the future development of their networks, a key component in the evolution to the gigabit NREN.[16] Some aspects of these plans are still unclear; for example, NSF has left to prospective bidders the choice of switching technology, from among those being investigated by the testbeds and elsewhere. However, the agency plans appear to be consistent with the target established by the testbeds. Initially, the agency networks will operate at lower bandwidths than the testbed networks, but they will incorporate more of the testbed technology as they evolve over time to meet the goal of the gigabit NREN.

The rate of NREN evolution is less dependent on technology issues than on delays in the selection of service providers.

Today, the highest bandwidth of the agency networks is 45 Mb/s; it appears that they will move to 155 Mb/s in 1994, with 622 Mb/s the highest rate that is realistically achievable by 1996. The rate of evolution is less dependent on technology issues than on delays in the process by which the Federal agencies select suppliers of NREN network services. Because agency choices of technologies and suppliers have broad implications for the Internet and the national information infrastructure in general, there have been several disputes over agency plans (see p. 7). While the NREN program has created a high level of interest in advanced networks, further delays in the deployment of agency networks may reduce the degree to which they will play the role of technology pioneers.

The agency networks' evolution depends in part on the timely deployment of the necessary high bandwidth transmission infrastructure by the telecommunications carriers. Computer networks generally use links supplied by the carriers—the network operators do not normally put their own fiber in the ground. The carriers' networks already have gigabit-capacity fiber installed, but today the capacity is usually divided among thousands of low-bandwidth channels used for telephone calls. New transmission equipment, the electronics at each end of the fiber, is required to allow the fiber's capacity to be divided into the high-bandwidth channels needed by the gigabit NREN. This equipment is being used in the testbeds and is becoming available commercially, but is very expensive.

The testbed applications research helps researchers to understand how the NREN would be used to achieve the science goals of the overall HPCC program. For example, some of the testbed applications show how networks can be used to bring greater computer power to bear on complex simulations such as the Grand Challenge problems. They may also show how networks can be used to help researchers collaborate—the Grand Challenge teams are expected to involve scientists at widely separated locations. In 1992, the NSF supercomputer centers proposed the concept of a "metacenter," which uses a high-speed network to link the computing power of the four NSF supercomputer centers.

The testbeds do not address all of the technology issues that are key to the future development of the NREN. Because the NREN will develop from the federally funded segment of the current Internet, it is affected by issues related to the growing number of users of the Internet. This growth in the number of users is straining some of the Internet protocols, and their future development is a topic of intensive study and debate

[16] NSF, op. cit., footnote 8; Leighton, op. cit., footnote 8.

within the Internet community. Also, the testbeds are not looking at applications that would be used by a broad range of users in the near term, or at issues related to making the Internet applications easier to use.

OTHER NETWORKS

One of the roles of the NREN is to serve as a testbed in itself, demonstrating technology that will then be deployed more broadly in the national information infrastructure. The testbed program will also impact the evolution of the national information infrastructure more directly, bypassing the intermediate stage of deployment in the NREN. This is because the network technology used in the testbeds reflects near-term industry planning. While the testbeds have emphasized higher bandwidths and more specialized applications than are of immediate commercial importance, the testbed networks reflect ideas that figure prominently in industry plans and, wherever possible, use equipment that conforms to emerging standards.

For example, many of the testbeds use a switching technology called Asynchronous Transfer Mode or ATM. This technology has become central to telecommunications industry planning because it is designed to support many different kinds of services—today's telephone network switches are limited mainly to carrying ordinary telephone calls. ATM can support Internet-type services such as will be used in the NREN, and also video, voice, and other data communications services—the carriers plan to use ATM to enter a variety of markets. Although ATM has been widely accepted by the telecommunications industry and progress has been made towards its implementation, there are a number of unresolved research issues. The testbeds are providing a large-scale opportunity to test this technology and possibly provide input to the standards process.

2 The Internet

The gigabit National Research and Education Network (NREN) is to develop from the current Internet, a "network of networks" that connects users in all parts of the United States and around the world. The Internet allows users to communicate using electronic mail, to retrieve data stored in databases, and to access distant computers. The network began as an Advanced Research Projects Agency research project to investigate computer networking technology, and in slightly over 20 years has grown into an essential infrastructure for research and education. The NREN initiative and associated research programs are intended to support the further evolution of research and education networking, broadening access to the network and enabling new applications through the deployment of advanced technologies.

Federal support to further the development of networks that support research and education communications is directed primarily at upgrading the Federal "backbone" networks that have formed the core of the Internet.[1] These networks include the National Science Foundation's NSFNET backbone, the NASA Science Internet (NSI) (figure 2-1), the Department of Energy's Energy Sciences Network (ESnet), and the Department of Defense's DARTnet and Terrestrial Wideband Network (TWBnet). The NASA and DOE networks are primarily intended for traffic related to the mission of the supporting agency, while the current NSFNET backbone serves users in a broader range of disciplines in universities, supercomputer centers, and industry research laboratories. The DOD networks support research and development of new communications technologies. The Federal

Federal agency networks will form the core of the gigabit NREN.

[1] Office of Science and Technology Policy (OSTP), "Grand Challenges 1993: High Performance Computing and Communications," p. 18.

Figure 2-1—Federal Agency Backbone Network

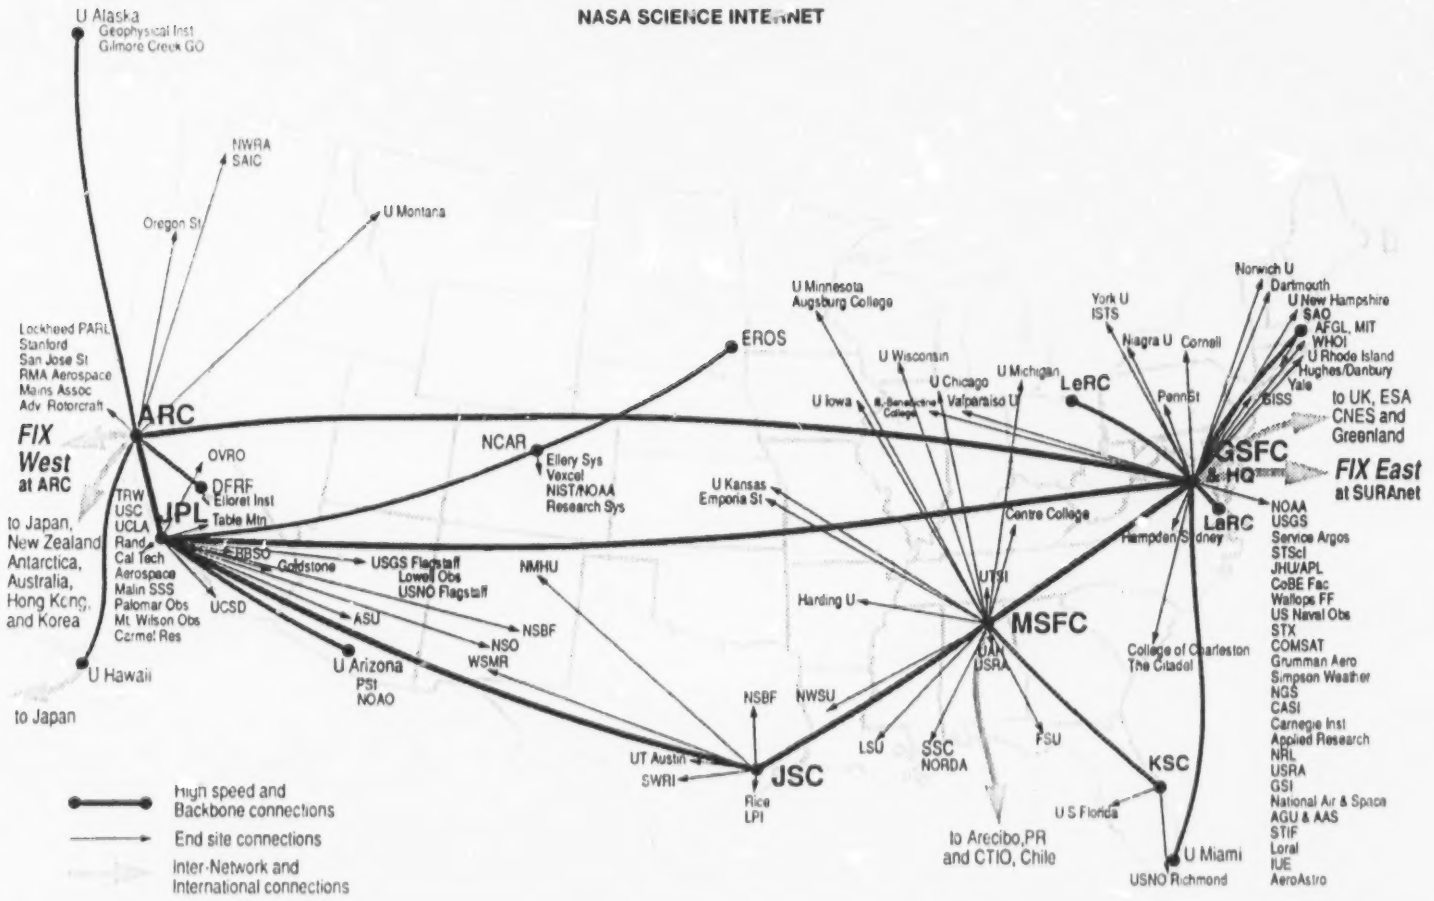

Important sites on the NASA Science Internet (NSI) include Ames Research Center (shown as "ARC"), Jet Propulsion Laboratory (JPL), Johnson Space Center (JSC), Marshall Space Flight Center (MSFC), and Goddard Space Flight Center (GSFC). NSI is connected to other Federal networks at the Federal Internet Exchanges (FIXes) East and West.

SOURCE: National Aeronautics and Space Administration (NASA).

Figure 2-2—Regional Network

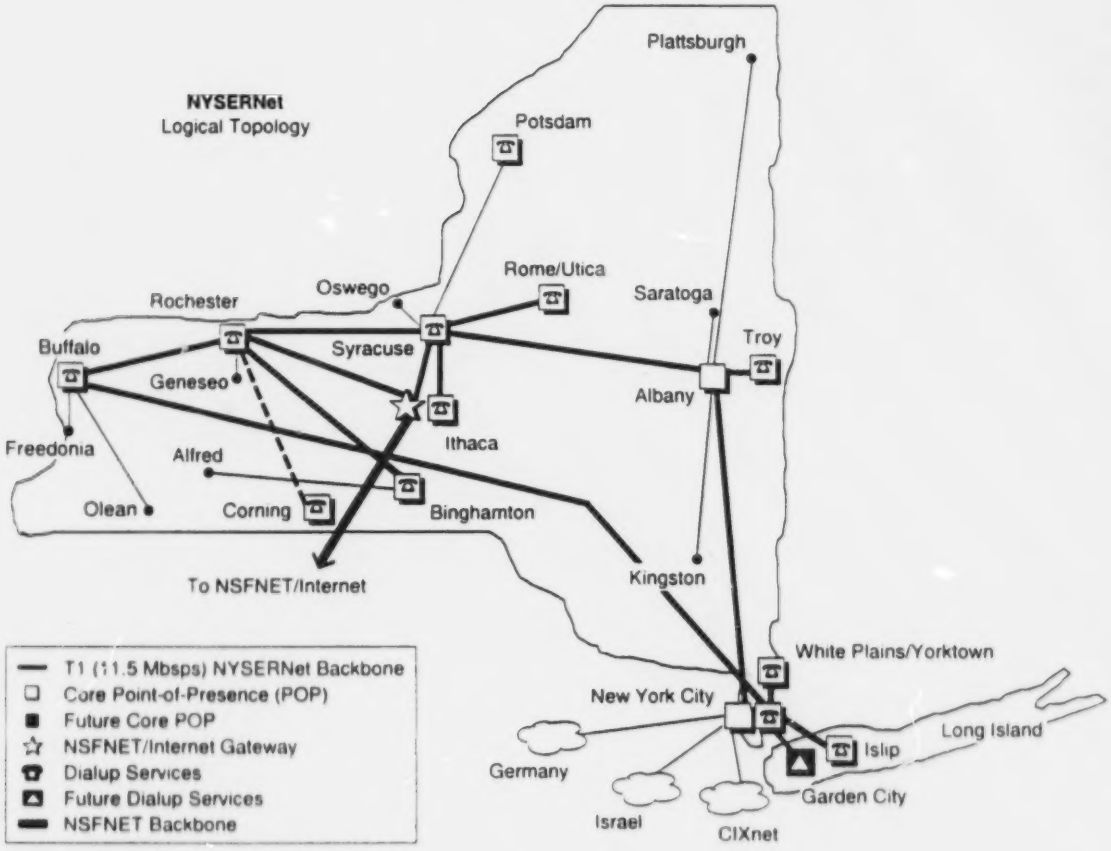

SOURCE: NYSERNet.

networks are interconnected at FIXes (Federal Internet Exchanges) at NASA's Ames Research Center in California and at the University of Maryland.

Upgrading the agency-supported backbones is not the only thing needed to improve research and education networking. The majority of users in universities, schools, and libraries do not have direct access to one of the backbone networks. These users rely on thousands of other networks that, together with the Federal agency backbones, form the Internet. These networks are interconnected, and information typically travels through several networks on its way from one user to another. In order to provide good performance end-to-end, all of the Internet's networks will need to evolve in a coordinated fashion, matched in capability and performance.

Most of the Internet's networks are "campus" or "corporate" networks, connecting users within a university or a company. Campus and corporate networks may in turn be interconnected by "regional" networks. For example, NYSERNet (New York State Education and Research Network) connects campuses and industrial customers in New York State (figure 2-2) and BARRNET (Bay Area Regional Research Network) does the same in northern California.

Figure 2-3—The NSFNET Hierarchical Structure

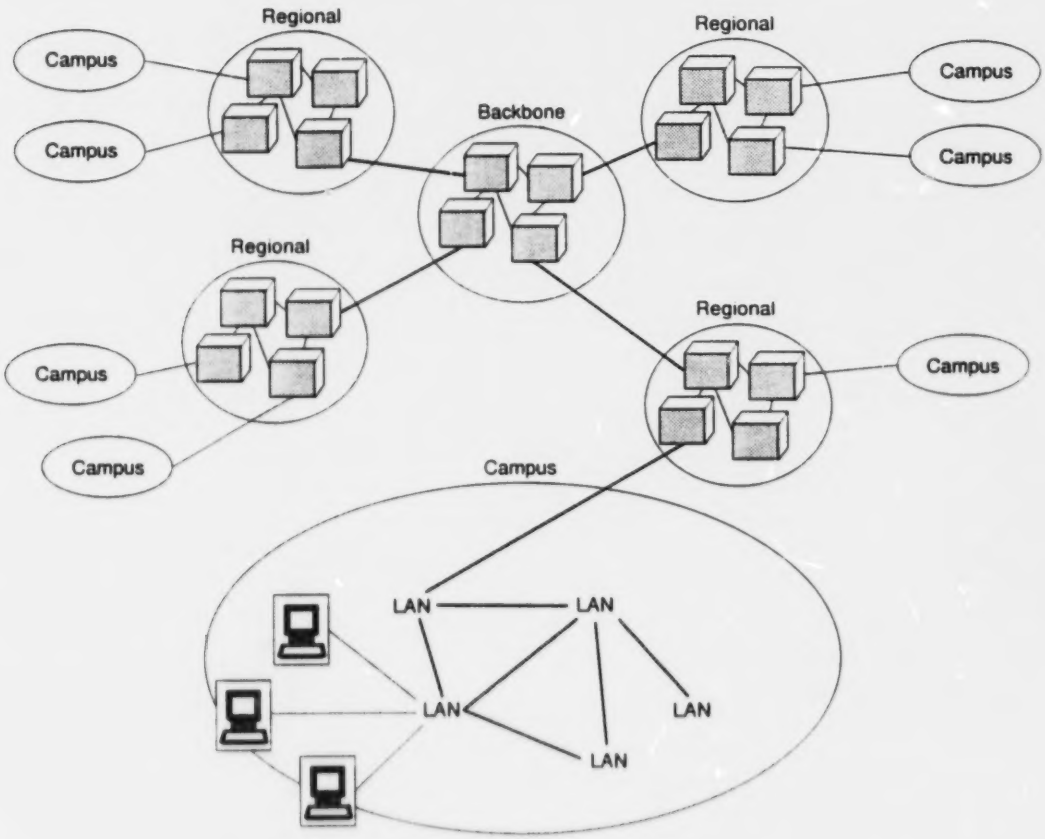

SOURCE: National Science Foundation (NSF).

Regional networks also provide a connection between campus networks and the national NSFNET backbone that carries traffic to other regions.[2] The regional networks, and the resulting three-tier structure of campus, regional, and backbone networks (figure 2-3), evolved with support from the National Science Foundation.[3]

The Internet also includes several networks that provide service on a for-profit basis.[4] The government investment in developing and demonstrating Internet technology during the 1970s and 1980s has created opportunities for the private sector to sell Internet services. The effectiveness of the Internet technology has been proven, and a

[2] NASA and DOE sites are connected directly to the agency networks. However, NASA and DOE rely on the regional networks and the NSFNET backbone to connect to university researchers participating in NASA and DOE projects.

[3] For a description of evolution of the regional networks and the three-tier structure, see Richard A. Mandelbaum and Paulette A. Mandelbaum, "The Strategic Future of the Mid-Level Networks," Brian Kahin (ed.), *Building Information Infrastructure* (New York, NY: McGraw Hill Primis, 1992).

[4] Eric Arnum, "The Internet Dilemma: Freeway or Tollway," *Business Communications Review*, December 1992, vol. 22, No. 12, p. 31.

growing number of companies are now using the Internet to conduct business. Even though the NREN program continues government funding for the agency backbone networks, in order to upgrade them to gigabit speeds, government support is becoming less central to the Internet as a whole. New commercial providers of nationwide Internet services have emerged. In addition, NSF has been reducing subsidies to the regional networks, which are increasingly being asked to recover costs from users.

The availability of commercial services is leading to a change in the makeup of the users of the Internet. Until recently, corporate use of the Internet was restricted to scientists and engineers in research laboratories or engineering departments. In part, this was due to the history of the Internet as an experimental network. The limited use of the Internet by the private sector was also due to an "Acceptable Use" policy that reserved the federally supported backbones for research and education traffic.[5] The new commercial providers have no traffic restrictions, allowing the Internet to serve a wider range of users. Today's Internet users can have different security requirements[6], their technical sophistication varies, and the demands they place on the network's capacity differs.

One of the goals of the NREN program is to continue the trend towards provision of Internet services on a commercial basis, rather than solely as the result of a government subsidy.[7] The NREN program continues government support for networking, but the emergence of commercial providers is leading to changes in the mechanisms by which this support is provided. NSI and ESnet will continue to support agency missions, but the next-generation NSFNET backbone will be considerably different from the current NSFNET backbone. As part of its NREN plans, NSF has decided that much of the traffic that is currently carried by its NSFNET backbone will in the future be handled by commercial providers, encouraging the further development of this segment of the Internet.

The next-generation NSFNET backbone will support a narrower range of users and serve fewer sites. Today's NSFNET backbone serves many sites nationwide, connecting regional networks and supercomputer centers (figure 2-4). It is a "general purpose" backbone, carrying traffic ranging from ordinary electronic mail to advanced supercomputer applications. In the future, the backbone will primarily be used by the NSF supercomputer centers, in Ithaca, New York, Pittsburgh, Pennsylvania, San Diego, California, and Champaign, Illinois.[8] Other users, with more routine applications, will use services available from commercial providers. Without the current national backbone, the regional networks will have to make new arrangements for their interconnection (see ch. 5, p. 67).

The next-generation NSFNET backbone will continue to contribute to the objective of developing advanced network technology. The new backbone, together with the next-generation NSI and ESNET, will be one of the first networks to use the technologies studied by the gigabit testbeds described in chapter 4. The Federal networks will provide "experimental" services, not yet available from commercial providers. They will demonstrate and test new network

[5] For issues related to NSF's Acceptable Use Policy, see Hearings before the House Subcommittee on Science, Space, and Technology, Mar. 12, 1992, Serial No. 120.

[6] Gary H. Anthes, "Internet Security Risks," *Computerworld*, vol. 26, No. 48, Nov. 30, 1992, p. 55.

[7] "[T]he NREN Program has a series of synergistic goals [including] stimulating the availability, at a reasonable cost, of the required services from the private sector," Office of Science and Technology Policy (OSTP), "The National Research and Education Network Program: A Report to Congress," December 1992, p. 2.

[8] For a description of the NSF supercomputer centers, see U.S. Congress, Office of Technology Assessment, *High Performance Computing and Networking for Science*, OTA-BP-CIT-59 (Washington, DC: U.S. Government Printing Office, September 1989), pp. 9-10.

Figure 2-4—NSFNET Backbone

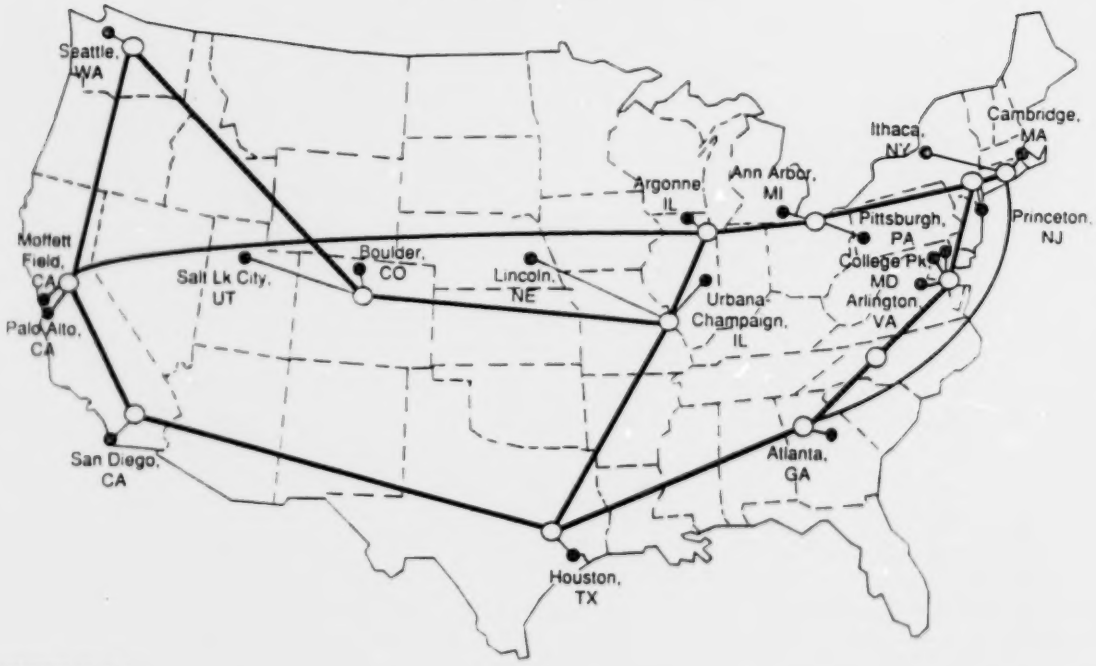

SOURCE: Merit, Inc.

technologies and applications before they are deployed more widely by operators of commercial networks.

Federal agencies may subsidize access to network services for users not at one of the backbone sites.[9] Today, NSF lowers the cost of networking for many users by directly subsidizing a general-purpose backbone and by providing subsidies to the regional networks. This strategy has contributed to broadening network access beyond major universities and supercomputer centers, to include many colleges and schools. In the future, many of these users will no longer be able to use the subsidized NSF network. Instead, Federal agencies may subsidize users' purchases of services from the commercial providers.

The NREN can then be viewed as many interconnected networks, developing from components of the current Internet. Some networks—the agency backbones—will be funded directly by the government. This part of the NREN is sometimes referred to as the "Interagency Interim NREN" or "NREN proper," and will use advanced network technologies to support high-end users, agency missions, and the science objectives of the HPCC program. Other Internet networks—such as existing regional networks or new commercial providers—may also carry NREN traffic, from users subsidized by the government, but would carry commercial traffic as well. These networks will likely use less sophisticated network technology than the agency backbone networks.

[9] "Over time, NSF will target its funding to those campuses which have financial impediments to connecting into the U.S. Internet." Robert Aiken et al., "NSF Implementation Plan for Interagency Interim NREN," May 1992, p. 4; "Federal funds. . .will also support users that serve Federal missions whether or not they access NREN through the agency networks," OSTP, op. cit., footnote 7, p. 3.

This two-part strategy—agency operation of advanced networks combined with subsidies for Internet access for certain groups of end users—represents a more detailed framework than the general NREN concepts and goals outlined in the High Performance Computing Act of 1991. It is expected to form the basis of NSF's forthcoming solicitation for the operation of its component of the NREN. It is also outlined in recently introduced legislation, the High Performance Computing and High Speed Networking Applications Act of 1993 (H.R. 1757), which would amend the High Performance Computing Act of 1991. However, there is concern in parts of the user community most affected by the change to an environment in which there is no longer a general purpose government operated network about the cost of commercial services and about the timing and management of the transition.

The remainder of this chapter describes the technology used in the current Internet. Chapter 3 provides an overview of emerging concepts that address some of the limitations of current network technology and might be used to construct gigabit networks. Chapter 4 describes the gigabit testbeds, NSF- and ARPA-funded prototype networks that are investigating these new technologies. Chapter 5 outlines NSF, NASA, and DOE plans for the deployment of the testbed technologies in their networks.

APPLICATIONS

From the users' perspective, an "application" is a task that the combination of the computer and the network enables them to perform. For example, a science teacher might use the Internet to locate information that can be used in a class, such as images stored in NASA databases, or databases containing tailored educational materials. Researchers use the Internet to track developments in their field, by exchanging information or drafts of papers and collaborating with other scientists.[10] In the business world, networks are increasingly used to track inventory or manage activities throughout a large company. In the future, networks may be used to help provide medical services to distant locations.

From a network engineering perspective, an "application" is a computer program that builds on the basic network service to allow a user to perform tasks. The application program provides interaction with the user; it does not handle the details of moving a message through the network to its destination. These functions are performed by communications software—a second program running on the computer—and specialized hardware that converts the computer's digital data to the format used by the network. When an applications program wants to send information to another computer, it hands the message to the communications software, which then formats the message and sends it over the network.

There are four major Internet applications—electronic mail (e-mail), file transfer, remote login, and news. Electronic mail is used to send messages to other users of the Internet, and for most users it is probably the application they use the most frequently. File transfer (File Transfer Protocol or FTP) is used to retrieve a "file" from another computer; a file could be a computer program, an article, or information from a commercial database. "Remote login" (Telnet) is used to control a distant computer; this is the application used to access a supercomputer or one of the other specialized computing resources on the Internet. "News" is a kind of bulletin board or discussion group—thousands of "newsgroups" address a wide range of different topics.

The current Internet applications are difficult to use. For example, it is difficult to find information resources on the network. First, the user has to know that the information exists somewhere reachable on the network, then where to find it, and, having found the database, how to locate the information in the database. A number of new

[10] For an overview of the wide range of uses for the Internet, see Daniel P. Dern, "Applying the Internet," *Byte*, February 1992, p. 111.

applications assist this process by acting as indexes or catalogues. Second, the user interface for most applications is often difficult to use, requiring a user to recall obscure commands. The difficulty in use is partly due to the Internet's heritage as an experimental network used mainly by scientists and engineers who were comfortable with arcane computer languages.

The existing Internet applications programs are beginning to be replaced by more sophisticated versions.[11] Today, for example, the Internet file transfer program, FTP, is used to retrieve a file from a distant computer, but a different program is used to retrieve a file stored on the "home" computer. Newer versions of these applications are "transparent," so that the user will not know whether a file is located on a distant computer, or that a program is executed on a different machine. These new applications are the beginnings of a foundation for "distributed computing," in which the computers on a network form an integrated system that performs as a single computer.

Applications and Network Technology

Some limitations of current applications are due to the applications software itself, but other limitations are due to the underlying network technology. One problem with current network technology is a shortage of bandwidth. Bandwidth is a measure of the amount of data that can be moved through the network in a given period of time, and is typically specified in terms of "bits per second." Because of the limited capacity of today's network, it is often impractical to move large amounts of data across the network—examples of large files are images (see box 2-A) and the data sets used in supercomputer applications.

A second limitation of current Internet technology is that it is best suited for applications that handle text or numerical data. The Internet is less effective when supporting applications that make use of "real-time" media such as video and sound.[12] In the case of video, this is due in part to the bandwidth limitation—high-quality video needs to move large amounts of data, and the necessary bandwidth is not available throughout the Internet. Support for video and sound is also limited because the performance of the Internet is highly variable. Because video creates the illusion of motion by sending a "stream" of pictures at regular intervals, a longer delay in the time it takes one of the pictures to get through the network interrupts the video information that is being displayed on the user's computer.[13] A new technology called "fast packet switching," discussed in detail in chapter 3, may provide the more consistent network performance that video applications need. Digital transmission and high bandwidth alone are not always sufficient to enable a network to carry video.

The limited capacity of the current Internet and the variability of its performance also constrain the use of sophisticated "distributed computing" applications. In distributed computing, one is able to treat the computers on a network as a single, more powerful computer. For example, two computers, exchanging data through the network as necessary, might be able to complete a computation in half the time needed by one computer working alone. If data takes too long to travel between the computers, however, the advantages of dividing a computation among several computers are lost. In the current Internet, the local area network (LAN) technology used in campus networks often performs better than wide area network (WAN) technology used in the

[11] For example, "distributed file systems" are beginning to replace the traditional File Transfer Protocol (FTP) application.

[12] Jeffrey Schwartz, "A Push for Packet Video," *CommunicationsWeek*, Aug. 3, 1992, p. 1.

[13] This problem is being attacked in a number of ways. New network architectures, described in chapter 3, try to reduce the degree of variation in network performance. Other researchers are investigating mechanisms that would compensate for the variable performance. For example, the receiving computer could "even out" some of the variation before the data is displayed to the user.

Box 2-A—Images and Video

Images

The screen of a computer's display is made up of many individual picture elements or "pixels," like the little dots that can be seen on television screens. By displaying each pixel with a different shade and different color, the computer forms an image on the screen. The greater the density of pixels, the higher the "resolution" of the image. The displays used for ordinary desktop computers usually have a few hundred pixels in both the horizontal and vertical directions, while a high-definition television display would have about 1,000 pixels vertically and about 2,000 horizontally. Even higher resolution displays are being developed for specialized medical, publishing, and defense-related applications.

The use of high-resolution images places considerable demands on computers and networks. Typically, each pixel on a screen is represented by 24 bits. A high-resolution display with 2,000 pixels horizontally and 2,000 pixels vertically has 4 million pixels (2,000 x 2,000 = 4,000,000). This means that 96 million bits are needed to represent the image (4 million x 24 = 96 million).

In the telephone network, voice conversations are sent through links that transmit 64 thousand bits per second. Using these links, an image represented by 96 million bits would take 25 minutes to send through the network. By contrast, it would take less than one-tenth of a second to send the same image through a gigabit network.

Video

Video is a series of images, sent many times a second at regular intervals in order to create the illusion of motion. Typically, 30 or 60 images are sent every second. In a low bandwidth network, in order to send this many images every second, the images have to be of very low resolution.

Two strategies have been adopted for accommodating image and video transport in networks. The first is to use *compression* techniques that reduce the number of bits needed for each image. Often, some parts of a scene do not have to be shown in great detail. Compression schemes for videotelephones sometimes rely on the fact that users are only interested in the "talking head," not the background. Sometimes little changes from one image to the next (if there is no movement in the scene), in which case the image data does not need to be sent again.

These techniques are being applied to the new high-definition television systems that are being studied by the Federal Communications Commission for selection as a U.S. standard. An uncompressed high-definition television signal that sends 30 images or "frames" every second, with a resolution of 1,000 pixels vertically and 2,000 pixels horizontally, needs about 1.5 gigabits per second. By contrast, new compression algorithms support high-definition television at bandwidths of 30 Mb/s or less, one-fiftieth the bandwidth required for the uncompressed signal.

The second strategy for accommodating video or images is to increase network capacity. Fiber optic technology can transport many more bits every second than the "twisted pair" copper wires that are used for today's telephone service. This background paper outlines some of the research being done on very high-capacity networks that can carry high-resolution video and images. However, even a "gigabit network" is not sufficient for certain kinds of very high-resolution video, and compression techniques might still be used.

SOURCES: Peng H. Ang et al., "Video Compression Makes Big Gains," *IEEE Spectrum*, vol. 28, No. 10, October 1991, pp. 16-19; Bernard Cole, "The Technology Framework," *IEEE Spectrum*, vol. 30, No. 3, March 1993, pp. 32-39; J. Bryan Lyles and Daniel C. Swinehart, "The Emerging Gigabit Environment and the Role of Local ATM," *IEEE Communications*, vol. 30, No. 4, April 1992, pp. 52-58.

Figure 2-5—Growth in NSFNET Traffic

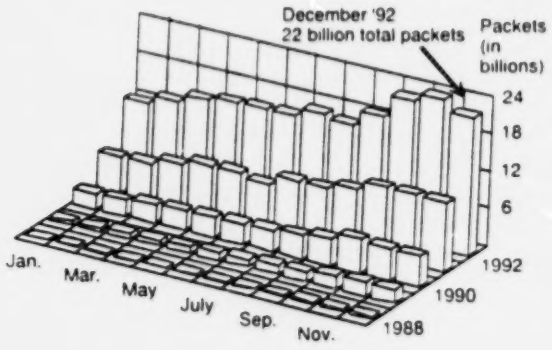

SOURCE: Merit, Inc.

regional and national backbone networks and thus distributed computing applications are used more widely in the local environment.

Limitations in network performance are becoming more apparent as computer technology advances. First, advances in computer power have resulted in demands for more bandwidth (figure 2-5). The size of the files that users would like to send through the network is increasing as a result of greater processing power and larger memories.[14] Some of the new "massively parallel" computers being studied as part of the HPCC Program may accelerate this trend (see box 2-B). Furthermore, the declining cost of computing power has allowed more users to connect to the network, creating more demand for the limited amount of capacity.[15] Second, computers are increasingly equipped with display technology that supports video-based applications. As video and sound begin to be processed by computers, there will be greater demand for networks that support this stream-type traffic. Today's networks were designed for an environment in which computers were restricted to working with ordinary text and numerical data.

In response to the limitations of today's networks and the trends in computer design, there is now a general vision of the type of services that future computer networks will have to support—larger, possibly image-oriented files, greater use of stream-type services such as video and sound, and more distributed computing. However, there are a number of issues that must be solved, and researchers are trying to learn more about the applications that users will need in the future. Because most network technologies support some types of applications better than others, arguments in the technical community about the best way to build broadband networks can often be traced to different assumptions about the expected mix of applications. One of the objectives of the NSF/ARPA gigabit testbeds discussed in chapter 4 is to learn about applications for advanced networks by encouraging collaboration among applications developers and network engineers.

PROTOCOLS

The Internet is a "packet-switched" network—a very different design from that used by the telephone network.[16] Data travels through the network as a "packet," a block of digital data consisting of the application's data and some extra information added by the communications software and hardware.[17] This information is sent either before the applications data in a "header," or after the data, in a "trailer," and tells the network the packet's destination address or instructs the receiving computer as to what to do with the applications data in the packet (figure 2-6). For example, the sending computer could

[14] Tim Studt, "Can High-Performance Networks Meet Future R&D Needs?" *R&D Magazine*, October 1992, pp. 30-34.

[15] Traffic on the NSFNET backbone is growing at a rate of 10% per month. OSTP, op. cit., footnote 7, p. 31.

[16] Vinton G. Cerf, "Networks," *Scientific American*, vol. 265, No. 3, September 1991, pp. 72-81.

[17] Often, packets are compared to the envelopes used in ordinary mail service—the extra information in the packet performs much the same function as the address on an envelope.

Box 2-B—Massively Parallel Computers

The conventional computers found on most desktops use a single processor. Programs for these computers consist of a list of instructions, to be executed one after another by the processor. Parallel computers are based on the idea that a computer with several processors can solve a problem more quickly than a computer with a single processor. Much of the HPCC Program's supercomputer design research focuses on the development of "massively parallel" computers with thousands of processors.

Supercomputers are expensive, high-performance machines that have been used mainly for numerical simulations in science and engineering. The first commercially important supercomputer, the CRAY-1, was first sold in 1976. It used a single processor, and achieved its high performance by careful attention to processor design and the use of specialized electronics. Over the next decade, supercomputer designers followed this basic model, trying to achieve the highest possible performance with a single processor.

By the mid-1980s, however, it became increasingly difficult to squeeze better performance out of traditional supercomputer designs, even as more exotic technologies were applied to the task. As a result, supercomputer designers began trying a different route to improved performance—the use of several processors. One approach involved a relatively small number of traditional high-performance supercomputer processors. For example, in 1983, Cray shipped a supercomputer that used four processors to speed up performance.

By contrast, the massively parallel approach to supercomputer design uses hundreds or thousands of low-cost microprocessors (processors that fit on a single semiconductor chip). The greater the number of processors, the more powerful the computer. In many cases, the microprocessors are the same as those used in high-end workstations. The performance of microprocessors increases every year, creating the potential for even more powerful massively parallel supercomputers.

Supercomputer centers and Federal laboratories have purchased several massively parallel supercomputers and are exploring their use in a number of applications. A major challenge for users of massively parallel supercomputers lies in the area of software. Massively parallel computers have to be programmed in new ways, because programs can no longer be thought of as a simple list of instructions. New algorithms, efficient ways of solving numerical problems, will have to be developed. Research on algorithms and software tools that take advantage of the potential of massively parallel supercomputers is one focus of the HPCC program.

SOURCES: Glenn Zorpette, ed., "Special Report: Supercomputers," *IEEE Spectrum*, vol. 29, No. 9, September 1992; pp. 26-41; Office of Science and Technology Policy, "Grand Challenges 1993: High Performance Computing and Communications," 1992, pp. 13-17; Carl S. Ledbetter, "A Historical Perspective of Scientific Computing in Japan and the United States," *Supercomputing Review*, vol. 3, No. 12, December 1990, pp. 48-58.

put a short code in the header to tell the receiving computer that the data belongs to an electronic mail message—this allows the receiving computer to process the data appropriately after receiving the packet.

Once the packets have been formatted they are sent out of the computer and through the network's web of links and switches. Switches receive packets coming in on one link and send them out on the next link in the path to their destination (figure 2-7). When the packet arrives at a switch, the switch scans the destination address and determines which link the packet should transit next. The Internet packet switches or "routers" are special computers that have been provided with connections to a number of links and programmed to carry out the switching functions.

The software in the routers and the users' computers implement "protocols," the rules that determine the format of the packets and the actions taken by the routers and networked computers. The Internet protocols are often referred to as TCP/IP (the acronyms refer to the two

330-073 O - 93 - 3 QL 3

Figure 2-6—Packet

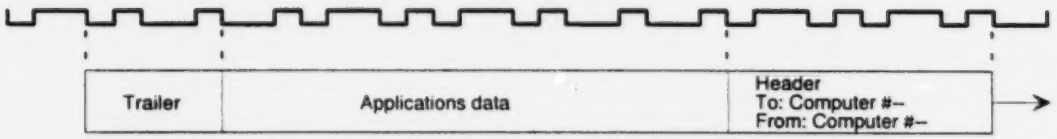

A packet is a block of digital data, consisting of data from the user's application and extra informationused by the network or receiving computer to process the packet. For example, the "header" might contain the "address" of the destination computer. A real packet would be several thousand bits long.

SOURCE: Office of Technology Assessment, 1993.

most important Internet protocols, the Transmission Control Protocol and the Internet Protocol.) Special protocols called "routing protocols" are used by the routers to keep a current map of the Internet and to determine the best path to a destination computer—for example, to choose a path that avoids heavily loaded networks.

One of the most important characteristics of the Internet is that the thousands of linked networks are independently operated; there is no central control of the Internet. However, by sharing the Internet protocols, the networks are able to exchange traffic. One of the functions of the Internet protocols is to mask differences in the technology used by the networks that make up the Internet. The campus networks' local area network technology differs from the wide area network technology used in the regional and national backbone networks, and there are many different local area network standards. The term "Internet" is short for "internetworking," the practice of linking technologically different and independently operated networks.

The future of the current Internet protocols is the subject of considerable debate in the Internet community. The most significant problem is that today's routing technologies are being strained by rapid growth in the number of connected networks and users.[18] The management of a complex and growing network has been one of the major challenges faced by the current NSFNET. A number of different proposals that would simplify the routers' task of finding paths through today's more complex Internet are being considered. The effect of increases in bandwidth on TCP/IP has also been debated in the technical community, and new protocols have been proposed. Many now believe that TCP/IP can continue to provide good service over gigabit networks, but internetworking in high bandwidth networks is a research topic in itself.

NETWORK COMPONENTS

A network is a complex system, consisting of many computer programs and hardware components such as links, computers, and switches. The overall performance of the network depends on how well these components work together. There are a number of potential bottlenecks—the rate at which data can be transferred from the computer's memory to the network, the rate at which data can be transmitted through the links, and the amount of time the switches need to decide where to send data next. Simply removing one of these bottlenecks does not guarantee that the overall performance of the network will improve. The emergence of fiber optics has removed the links as a bottleneck for the foreseeable future; the research projects described in chapter 4 show that this has exposed research issues in other parts of the network.

[18] Karen Lynch, "Internet Metamorphosis," *CommunicationsWeek International*, Aug. 10, 1992, p. 1.

Figure 2-7—Packet Switching

(a) Packet-switched communication

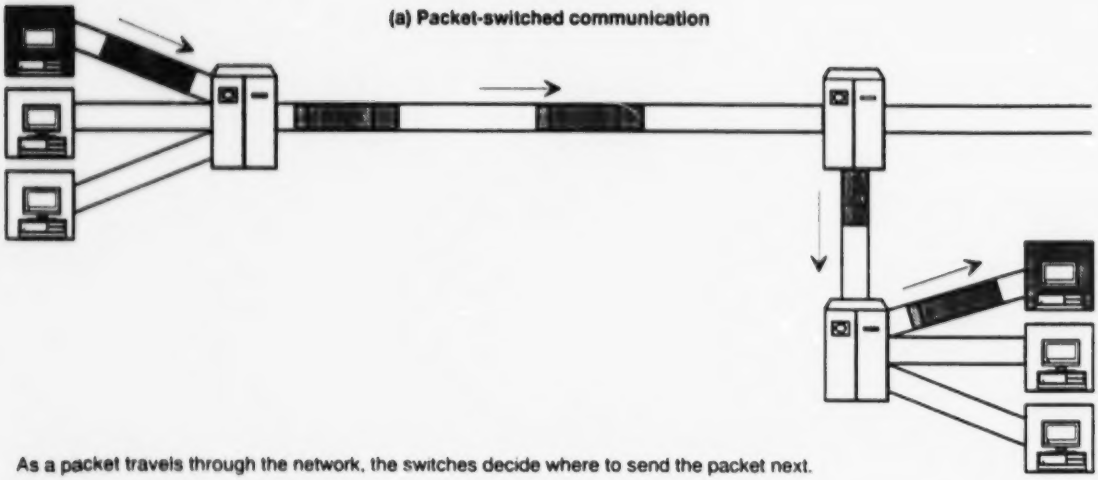

As a packet travels through the network, the switches decide where to send the packet next.

(b) Single link shared by several users

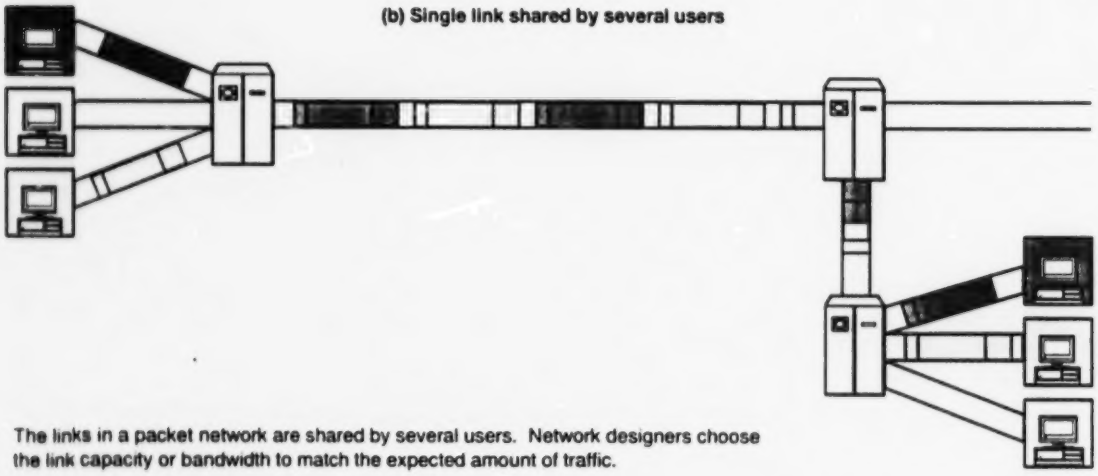

The links in a packet network are shared by several users. Network designers choose the link capacity or bandwidth to match the expected amount of traffic.

SOURCE: Office of Technology Assessment, 1993.

Computers

Many different kinds of computers are attached to the Internet, ranging from desktop personal computers costing a few hundred dollars to supercomputers that cost millions of dollars. Among scientists and engineers, the type of computer that is most widely used is the "workstation," a powerful desktop computer with enough processing power to support graphical user interfaces and high-resolution displays. For most of today's applications, almost any computer has enough processing power to attach to the Internet. The low bandwidth of the current Internet places few demands on computers for handling the communications functions, leaving much of the processing power free to run the applications.

One of the reasons for the creation of the NSFNET backbone was to provide access to NSF's four supercomputer centers. Recently, these supercomputer centers have begun to install "massively parallel" supercomputers. This new type of supercomputer attempts to achieve very high processing speeds by combining the processing power of thousands of smaller processors. Other supercomputers use a more traditional design, and are referred to as "vector" supercomputers. Each design may work best with certain kinds of computations; one of the objectives of the gigabit testbed research is to explore the use of networks to divide up problems in a way that takes advantage of the strengths of both vector and massively parallel supercomputers.[19]

Links

The digital links in computer networks usually use copper or fiber, but satellite and microwave links are also used. At each end of the copper or fiber is the transmission equipment, electronics that convert data into the optical or electrical signals that travel through the network. The capacity of the wires or strands of fiber depends on the characteristics of the material used and on the capabilities of the transmission equipment.

Today's Internet uses both low bandwidth links that operate over copper at a few thousand bits per second (kilobits per second or kb/s), and high bandwidth links that operate over fiber with a data rate of about 45 million bits per second (megabits per second or Mb/s). The test networks described in chapter 4 will use links that operate at a rate of one billion bits per second (a gigabit per second or Gb/s).

Typically, a single wire or strand of fiber carries many links at the same time. Through a process called "multiplexing," several low-bandwidth links can be aggregated into a higher bandwidth link. Gigabit-capacity fiber, for example, can be used either to carry several thousand low-bandwidth links used for telephone calls, or a single high-bandwidth link needed for a gigabit network.

The required link bandwidth depends on both the bandwidth requirement of each user and on the number of users sharing the link. One of the main reasons for upgrading the links in the NSFNET backbone from 1.5 Mb/s to 45 Mb/s in 1991 was to accommodate growth in the number of users. Growth in the use of routine applications can also be supported by simply adding more low-bandwidth links. However, new applications that need very large amounts of bandwidth to themselves require the deployment of higher bandwidth links. By increasing the link bandwidth to gigabit rates, the gigabit NREN will be able to support new classes of advanced applications, not just growth in the number of users.

Operators of wide area computer networks, such as the regional networks and the agency backbones, typically lease their links from the telephone companies. The telephone companies have already obtained the rights-of-way and have installed the transmission facilities for use in their core business, voice telephone service. Because of the reliance on telephone company facilities, discussions of computer network link bandwidth often use telecommunications industry designations of link capacity. For example, the current NSFNET backbone is often referred to as a "T3" network, after the industry designation of 45 Mb/s links. "T1" links, which operate at 1.5 Mb/s, are used in the current Department of Energy and NASA networks and in the regional networks. As the Federal networks are upgraded to bandwidths above the 45 Mb/s T3 rate, they will use a new family of transmission standards designed for high-capacity fiber optic links, called Synchronous Optical Network (SONET) (see table 2-1).

Universities and corporations install their own links in their buildings for use in local area networks. Local area networks can provide users with higher bandwidth than wide area networks—

[19] "A Union of Superpowers," *IEEE Spectrum*, vol. 28, No. 6, June 1991, p. 18.

Figure 2-8—Access Link

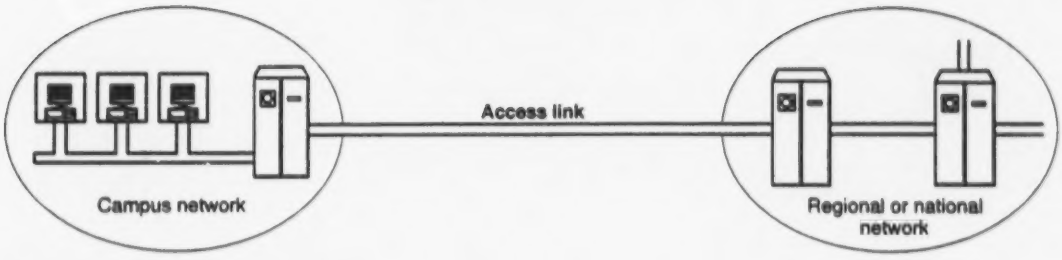

SOURCE: Office of Technology Assessment, 1993.

Table 2-1—Transmission Rates

Industry designation	Transmission rate
DS0	64 kb/s
T1	1.5 Mb/s
T3	45 Mb/s
SONET OC-3	155 Mb/s
SONET OC-12	622 Mb/s
SONET OC-48	2.4 Gb/s

SOURCE: Office of Technology Assessment, 1993.

this is due in large part to the high cost of high bandwidth wide area links. Because of the higher bandwidth available on local area networks, they have been used for experimentation with high-bandwidth distributed computing and video applications. In the future, however, users will want wide area networks that match the performance of local area networks; one of the objectives of the testbed project outlined in chapter 4 is to investigate high-speed wide area networking.

When campus networks arrange to be connected to the closest regional or national network, they obtain an "access" link (figure 2-8). This is usually leased from the telephone company, just as the links inside wide area networks are leased from the telephone company. The cost of the Internet service depends on the access bandwidth; high bandwidth access is extremely expensive. It is common to find local area networks operating at 10 Mb/s or 100 Mb/s, while the access link to the rest of the Internet operates at 56 kb/s or less (some organizations have 1.5 Mb/s access links, but these are considerably more expensive). Most individuals, schools, and small businesses are required to use their ordinary analog telephone line to access Internet services—a device called a "modem" is needed to send digital computer data over these lines, usually at 14.4 kb/s or less.

Switches

Packet switches in the Internet, also known as routers, direct packets to the next link in the path to their destination. Packet switched networks emerged to handle data communications, services not well supported by the "circuit switches" used for ordinary telephone calls (figure 2-9). Packet networks are more efficient for typical computer communications traffic—short transactions or "bursts" separated by periods of no traffic (box 2-C). In a packet network, several users share the same link—during the period in which one group of users is not using the link, other users can send their packets. In a circuit switched network, by contrast, each communication gets its own link. For this reason, circuit switches are most efficient when a communication involves a relatively long, steady stream of data such as video or voice.

While the Internet networks use telephone company links, the packet switches are usually not operated by the telephone companies. Instead, a second organization plans the network and installs the packet switches at the sites it has chosen—the involvement of the telephone company is usually limited to providing the links between the sites. From the perspective of the

Figure 2-9—Circuit Switching

(a) Telephone network

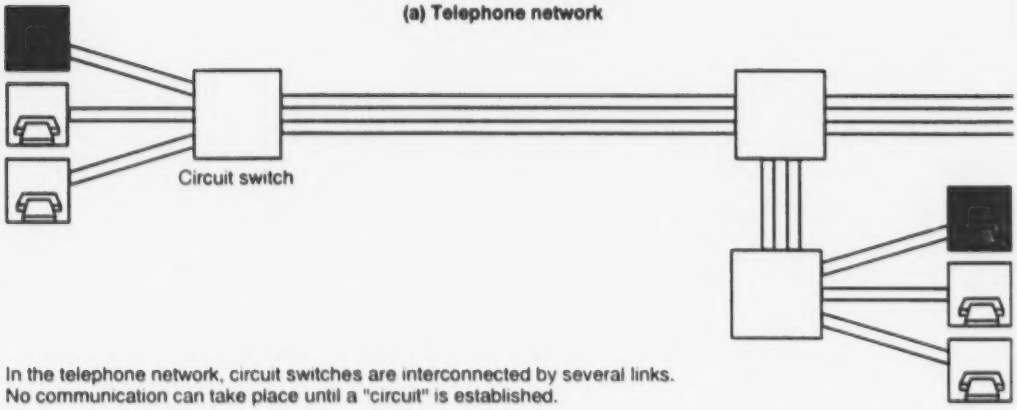

In the telephone network, circuit switches are interconnected by several links. No communication can take place until a "circuit" is established.

(b) Circuit-switched call

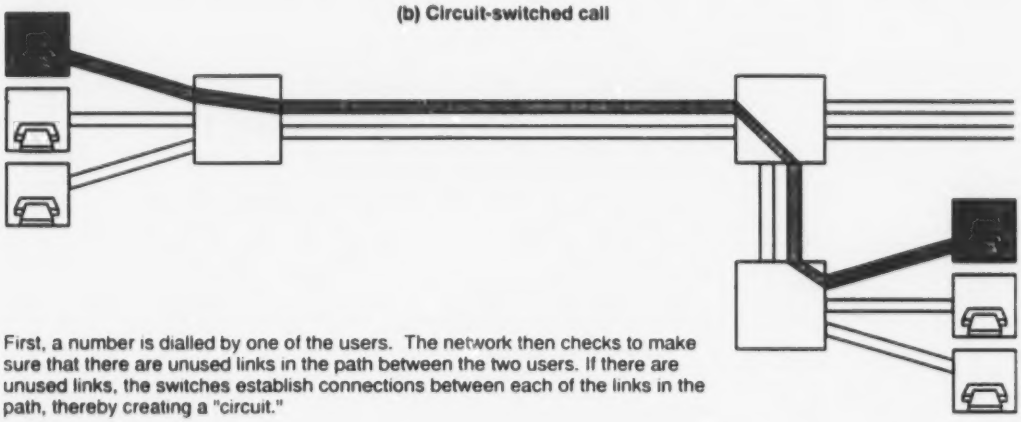

First, a number is dialled by one of the users. The network then checks to make sure that there are unused links in the path between the two users. If there are unused links, the switches establish connections between each of the links in the path, thereby creating a "circuit."

(c) Separate circuits for each call

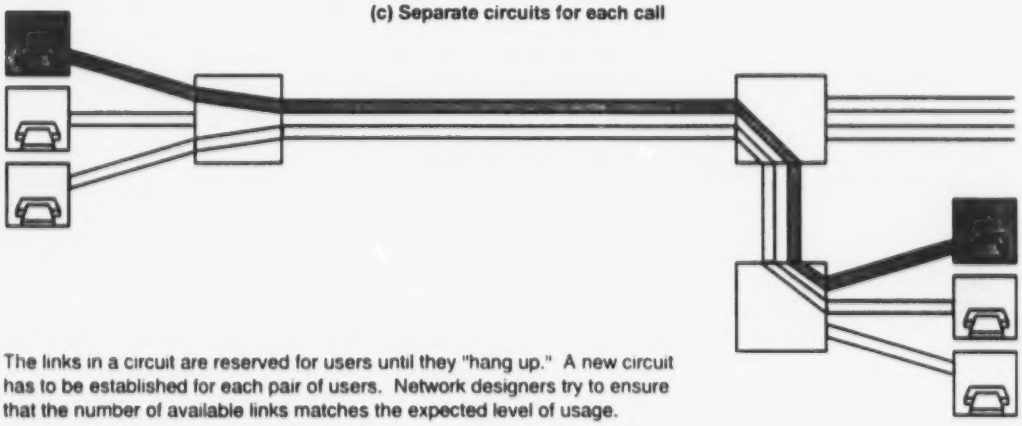

The links in a circuit are reserved for users until they "hang up." A new circuit has to be established for each pair of users. Network designers try to ensure that the number of available links matches the expected level of usage.

SOURCE: Office of Technology Assessment, 1993.

telephone company, the computer network traffic is just "bits" traveling over its links—the telephone company's equipment does not make decisions about where to send the packets. Beginning in the mid-1970s, the telephone companies began installing some packet switches in their networks in order to support the growing data communications market, but their efforts to enter this market were considered to be unsuccessful.

The processing power required of a packet switch depends on the link bandwidth and the complexity of the network. As the link bandwidth increases, switches must be able to process packets more quickly. The processing power needed will also increase as the network gets larger and more complex, because it becomes more difficult to determine the best path through the network. Currently, the NSFNET backbone's router technology does not allow the use of applications that need more than 22.5 Mb/s, half the potential maximum of a 45 Mb/s T3 network.[20] This shows how the overall performance of the network depends on many different components; increasing the link bandwidth is not the only requirement for an advanced network.

THE INTERNET AND THE PUBLIC SWITCHED NETWORK

In some ways, the Internet and the "public switched network" that is operated by common carrier telephone companies are separate. They differ in the services they provide—the telephone network mainly provides ordinary voice communications services, while the Internet provides data communications services such as electronic mail and access to remote computers. They also differ as to the communities that they serve—almost everyone has a telephone, while the Internet and other computer networks primarily serve users in the academic community or in industry. Finally, they differ in their network technology—the Internet is a packet-switched network, while the telephone network is a circuit-switched network. However, the Internet and the telephone network are related in a number of ways. Any discussion of the evolution of networking has to consider both the traditional telecommunications companies and the Internet community.

First, the Internet and the public switched network are related in that the links in wide-area computer networks are usually supplied by the telephone companies—computer network operators do not usually put their own fiber in the ground. As a result, the availability of new computer network capabilities can depend on the extent to which the telephone companies deploy advanced transmission facilities, and on the cost of leasing the links.

The availability of advanced transmission facilities varies[21], depending on whether a computer network will operate over the telephone network's "interoffice" or "local loop" segments. Most of the links required for a wide area network such as the NSFNET backbone operate over the interoffice core of the telephone network, which has largely been upgraded to optical fiber and digital transmission. The telephone companies upgraded this part of their networks in part to achieve operational savings, even when delivering existing services.

However, access links, such as those between a campus and a regional network, need to use local loop facilities. For the most part, this segment of the telephone network still consists of copper, analog lines. Large users are able to avoid this bottleneck by making special arrangements with the local exchange carrier for higher bandwidth digital lines. However, individuals, schools, and small businesses generally have to

[20] Alan Baratz, IBM, personal communication, Feb. 3, 1993.

[21] National Telecommunications and Information Administration, Department of Commerce, "Telecommunications in the Age of Information," October 1991, pp. 97-109.

Box 2-C—Packet Switching and Circuit Switching

Computer networks such as the Internet use packet switches, which direct packets from link to link through a network. Today's telephone network, by contrast, uses circuit switches. Each type of switching technology works best with different kinds of communications. Packet switching is more efficient for the transfer of typical computer communications traffic such as files of text or numerical data (figure 2-C-1). Circuit switching, on the other hand,

Figure 2-C-1—Packet Switching More Efficient for Data

(a) Data communications does not use circuits fully

Circuit switching can be used for computer communications. Here, circuits have been set up between two pairs of computers. However, computer communications often have a "bursty" character -- periods in which data is sent followed by periods of "silence." When no data is sent, the circuit's capacity goes unused. The capacity is used more efficiently when the communications involve a steady flow of information, such as video or voice transmission.

(b) Link sharing makes packet networks more efficient for data

In a packet-switched network, several users' traffic shares the same link. If one user is not using the link's capacity, it can be used by others. The figure shows bursts of data assembled into packets and travelling through the network on the same link. Here, one link's capacity is sufficient to handle communications between both pairs of users, freeing the second link for other uses.

SOURCE: Office of Technology Assessment, 1993.

use the combination of a modem and their telephone line to access computer networks. The bandwidth of such an arrangement is relatively low—only a few kb/s—and is clearly a bottleneck that limits widespread use of sophisticated services.

The telephone network and computer networks are also related in the sense that the traditional telecommunications providers are beginning to offer data communications services, including Internet services. In the past, efforts by the industry to enter this market have not been successful. This has been attributed to a "culture clash"—a lack of understanding of computer network technology and of the needs of users of computer networks. However, the telephone com-

can provide the consistent performance needed by video or voice traffic (figure 2-C-2). One of the objectives of the research described in chapter 3 and chapter 4 is to develop switches that combine the efficiency and flexibility of packet switching with the consistent performance of circuit switching.

Figure 2-C-2—Circuit Switching Better for Voice or Video

(a) Variable performance due to packet network link sharing

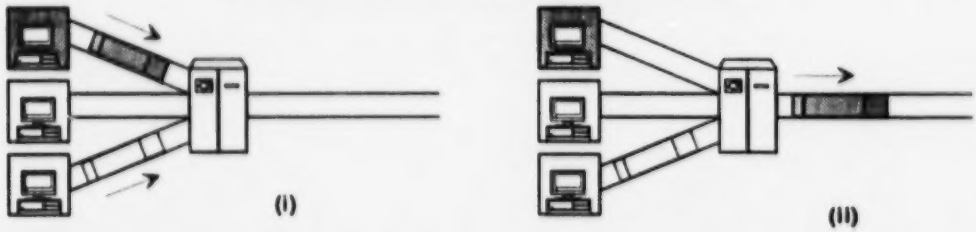

If two packets arrive at a switch at the same time and need to use the same outgoing link (i), one of the packets will have to wait (ii). It is difficult for a user to know in advance what the network performance will be. The packet may experience no delay (the dark gray packet), or it may have to wait at each switch (the light gray packet). This variation in delay has limited the use of packet networks for time-sensitive communications such as video or voice.

(b) Circuit switched performance is predictable

In a circuit-switched network, each communication has its own circuit. Users' information travels through the network without being affected by the characteristics of other communications (i)–(ii). The time needed for information to travel through the network will always be the same.

SOURCE: Office of Technology Assessment, 1993.

panies hope to play a more active role in this market.

The telephone companies have two main competitors in this venture. First, there are already a number of commercial providers of Internet services and other data communications services. These providers lease lines from the telephone companies, install packet switches, and operate their network without any further involvement from the telephone companies, sharing their network's capacity among different groups of users for a fee. The current T3 NSFNET backbone is obtained as a service from one of these commercial Internet providers.

Second, many users choose to operate "private networks"—they build a network of their own

using leased lines and bypass the public network. Most corporations use this strategy to interconnect local area networks at different sites within their organization. Equipment used in private networks is provided by computer companies and others, who have taken advantage of the telephone companies' lack of success in providing data communications services. United States firms that specialize in the development of routers and other equipment for private networks are world leaders and are among today's fastest growing companies.[22]

The telephone companies have introduced a number of new packet-switched services that are intended to encourage users to abandon their private networks.[23] One of these services is called Switched Multimegabit Data Service (SMDS); another is called Frame Relay.[24] The SMDS and Frame Relay switches do not understand the Internet protocols, but they can still be used to carry Internet traffic. The Internet packets are "encapsulated," or put inside an SMDS or Frame Relay "envelope," and sent through the network; at the other end the Internet packet is extracted and delivered to the computer. The carriers view SMDS and Frame Relay as transitional steps to a new technology called Asynchronous Transfer Mode (ATM), described in chapter 3. They can potentially be used to provide data communications services up to 45 Mb/s.

Because of the interrelationship between the Internet and the public switched network, the evolution of the Internet is affected by two different sets of standards committees. The telecommunications industry standards affect mainly low level issues, such as transmission standards, but some of the standards for new telecommunications industry packet switched services may play a role as well. The most important international standards group is the CCITT (International Telegraph and Telephone Consultative Committee). The CCITT is a technical committee of the International Telecommunications Union (ITU), a specialized agency of the United Nations that is headquartered in Geneva.[25] United States telecommunications standards are the responsibility of Committee T1, which is accredited by the American National Standards Institute (ANSI) and sponsored by the Exchange Carriers Standards Association (ECSA).[26] Telecommunications industry standards setting has often been criticized as excessively bureaucratic and slow.

By contrast, the Internet standards community, which addresses higher level issues related to routing, the TCP/IP protocols, and applications, is more informal. Much of the work is done by electronic mail, and there is a greater emphasis on proving that something works before it is standardized.[27] The two groups responsible for Internet standards are the Internet Engineering Task Force (IETF) and the Internet Activities Board (IAB). The IETF has a number of different working groups, each looking at a different aspect of the Internet's operation.

---

[22] G. Pascal Zachary, "U.S. High-Tech Firms Have Begun Staging Little-Noticed Revival," *Wall Street Journal*, Dec. 14, 1992, p. 1; G. Pascal Zachary and Stephen Kreider Yoder, "Computer Industry Divides Into Camps of Winners and Losers," Jan. 27, 1993, p. 1; Alan Deutschman, "America's Fastest-Growing Companies," *Fortune*, vol., 126, No. 7, Oct. 5, 1992, p. 58.

[23] Another market for these services is the smaller companies that cannot currently justify private networks.

[24] Tim Wilson, "Local Carriers Lay Out Data Service Agendas," *CommunicationsWeek*, May 25, 1992.

[25] G.A. Codding and A.M. Rutkowski, *The International Telecommunications Union in a Changing World* (Dedham, MA: Artech House, 1982).

[26] I.M. Lifchus, "Standards Committee T1—Telecommunications," *IEEE Communications*, vol. 23, No. 1, January 1985, pp. 34-37.

[27] Carl Malamud, "Stacks: Interoperability in Today's Computer Networks," (Englewood Cliffs, NJ: Prentice Hall, 1992), p. 223.

Broadband Network Technology 3

Advances in computer technology are driving the requirements for broadband networks. Because of increases in the processing power of computers, there is a need for higher bandwidth networks. Computers are increasingly able to execute "multimedia" applications, so it is expected that future networks must be able to carry several kinds of traffic. Broadband networks will lead to applications that are used for a wider range of problems, with more emphasis on image-based communications.

The computer and telecommunications industries have conceived broadband network designs for these requirements. Fiber optic links are a key component of these networks. However, replacing the smaller capacity links in current networks with higher bandwidth fiber optic links is not all that is needed: Improvements in protocol and switch design must also follow. Future switches will have more processing power, in order to keep pace with the faster flow of traffic through the links. They will also be designed in a way that allows them to handle different types of traffic. Today's switching technologies do not have this capability—packet switches only handle text and numerical data efficiently, the telephone network's circuit switches are best suited to voice traffic, and special networks are needed for video. The "integrated services" concept envisions networks that use the same links and switches for all types of traffic, instead of different technologies for video, data, and voice.

Broadband networks use new switch technologies.

BROADBAND APPLICATIONS

The new high bandwidth integrated services networks would improve the performance of existing applications and enable new applications. Existing applications, such as electronic mail or

databases, could be augmented through the use of image files and video clips; higher bandwidth networks would also allow the faster transfer of large files of supercomputer data. Support for real-time high-resolution video would expand possibilities further, allowing videoconferencing or the display of output from a scientific instrument, such as a telescope. More generally, the combination of more powerful computers and integrated services networks will permit wider use of two new categories of applications—multimedia applications and distributed computing.

Multimedia Applications

Multimedia applications take advantage of the capability of high-bandwidth integrated services networks to deliver many different kinds of data—video, image, audio, and text and numerical data. They also take advantage of the processing power of advanced workstations and other devices attached to the network, allowing users to edit, process, and select data arriving from a variety of sources over the network.[1] Multimedia applications have been suggested for a large number of areas[2], including education and health care. There are many different concepts for delivering multimedia services to the home, such as multimedia catalogues for home shopping, information services, entertainment video, and videotelephone services. Many segments of both service and manufacturing industries are increasingly using image-based applications—for example, computers are widely used in the publishing and advertising industries to compose pages using high-resolution images.

Multimedia is also the foundation for a new category of applications that use the combination of computing and communications to create a "collaborative" work environment in which users at a number of scattered sites are able to work together on the same project.[3] For example, an application might allow several researchers to work on the same set of experimental data at the same time—any processing done by one researcher would automatically be shown on the other researchers' displays. Videoconferencing and collaborative applications might allow closer interaction between researchers in different places. It is expected, for example, that the teams working on the Grand Challenges will include scientists at many locations.

For researchers, "visualization" provides a way to represent large amounts of data in a more understandable form; it uses images and video to show the results of simulations or experiments[4] (box 3-A). For example, the results of a simulation of a city's air quality could be shown as an image, with the concentration of a particular chemical indicated by different colors and color intensity. If a researcher wanted to review the evolution of the air quality over time, a series of images could be used to create a video segment showing the change in pollutant concentration. Other programs running on the workstation could be used to process the data further, perhaps by examining one part of an image more closely or by comparing the simulation data to experimental data.

In education, multimedia could be used in computer-based instructional materials. Multimedia databases would give students and teachers access to image and video data. Videoconferencing and collaborative applications could enable closer interaction between teachers and students at multiple locations. For example, it might

[1] Special Issue: Multimedia Communications, *IEEE Communications*, vol. 30, No. 5, May 1992.

[2] Michael L. Dertouzos, Director, MIT Laboratory for Computer Science, testimony at hearings before the Joint Economic Committee, June 12, 1992.

[3] Sara A. Bly et al., "Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment," *Communications of the ACM*, vol. 36, No. 1, January 1993.

[4] Matthew Arrott and Sara Latta, "Perspectives on Visualization," *IEEE Spectrum*, vol. 29, No. 9, pp. 61-65.

As part of the CASA testbed research described in chapter 4, a gigabit network will be used to combine data from a variety of sources, such as satellites and digital elevation models, to create three-dimensional views.

be possible to better emulate the classroom environment by allowing more two-way communication than is currently possible. Students might also be able to select a particular view of an experiment being demonstrated by a teacher. In health care, transfer of high-resolution images, such as x-ray and MRI data, combined with videoconferencing and other collaborative applications, could allow doctors to consult with specialists in other areas of the country.[5]

Distributed Computing

Other researchers have begun to consider the relationship between computers and communications in a more general way. "Distributed computing" uses the network to combine the processing power and memory of multiple computers. It is then possible, for example, to combine several low-cost workstations to achieve performance comparable to that of a supercomputer—a very expensive machine to purchase and operate. Computations can also draw on data stored in many different locations. Distributed computing becomes feasible as the network connecting the computers becomes less of a bottleneck, allowing them to work more closely together.

It may also be useful to do distributed supercomputing—using the network to provide proc-

[5] M. Niel Ransom and Dan R. Spears, "Applications of Public Gigabit Networks," IEEE Network, vol. 6, No. 2, March 1992, p. 31.

Box 3-A—Interactive Visualization Using Gigabit Networks

Scientists at the University of Wisconsin's Space Science and Engineering Center have developed a software package that allows "interactive" visualization of data computed by a model. This photograph of a workstation's display shows both a computed forecast of cold fronts moving across the North Atlantic and the "control panel" that the scientist can use to control the images displayed.

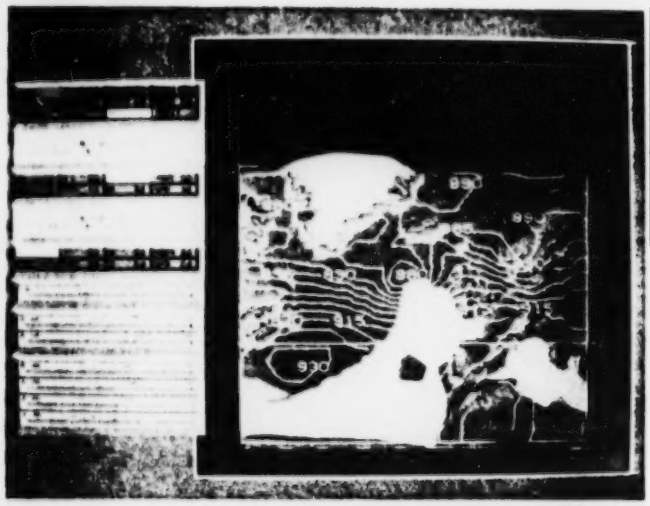

Bill Hibbard, University of Wisconsin.

Computers generate new kinds of images, that change in response to their users' needs. If the scientist selects the command "animate," in the upper left hand corner of the control panel, the workstation will display a succession of images that show the evolution of the storm over time. Other commands allow the user to rotate and "zoom" the images, to look at them from any angle.

The time required for a workstation to compute a new image, in response to a user command, can be significant. A supercomputer would be able to reduce the response time and allow interactive exploration of the data computed by the model. However, like many research institutions, the University of Wisconsin does not have a supercomputer.

The Internet could be used to send data to one of the NSF's supercomputer centers. However, the data rate of today's Internet is too low—the advantages of speeding up the computation by using a supercomputer are outweighed by the time needed to transfer the model data to and from the supercomputer. With a gigabit network, the communication time would no longer be a bottleneck.

As part of the BLANCA testbed's applications research (see ch. 4, p. 56), University of Wisconsin scientists will use a gigabit network to support interactive visualization of large data sets. The user's commands would be sent from the workstation through the network to a supercomputer at the National Center for Supercomputing Applications, in Champaign, Illinois, which would do most of the visualization processing and send the image data back to the workstation for display. This testbed research may serve to demonstrate a way for the majority of research institutions that do not have supercomputers to do interactive visualization.

SOURCES: William Hibbard, University of Wisconsin—Space Science and Engineering Center, Mar. 11, 1993; William Hibbard, David Santek, and Gregory Tripoli, "Interactive Atmospheric Data Access Via High-Speed Networks," *Computer Networks and ISDN Systems*, vol. 22, pp. 103-109.

essing power that exceeds that of a single supercomputer[6] (see box 3-B). Supercomputer applications are often simulations of real-world phenomena; for example, airplane designers do not need to build a model of an airplane and test it in a wind tunnel, but can simulate the flow of air around the airplane on a supercomputer. Unfortunately, for many interesting problems the process-

[6] Charles E. Catlett, "In Search of Gigabit Applications," *IEEE Communications*, vol. 30, No. 4, April 1992, pp. 43-45.

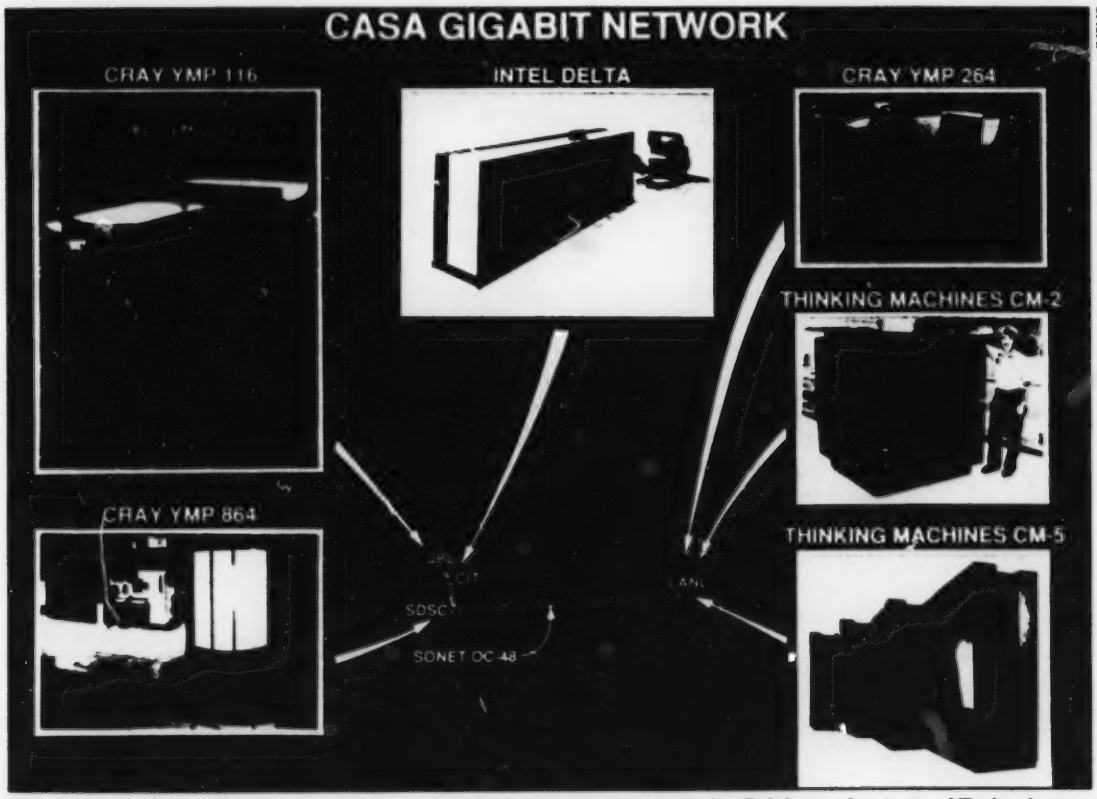

The CASA testbed will use a gigabit network to link supercomputers at the California Institute of Technology (CIT), Jet Propulsion Laboratory (JPL), Los Alamos National Laboratory (LANL), and the San Diego Supercomputer Center (SDSC).

ing time required with even the fastest supercomputer can be significant.[7] Researchers hope to reduce this time by using multiple computers in parallel and linking them through a network. The network could also connect the supercomputers to scientific instruments or massive remote databases that would provide data to be used in the calculations.[8]

Applications and Network Requirements

Two requirements will be placed on future networks. First, they will need to have much higher bandwidth than today's networks, in order to keep pace with advances in computer technology and support bandwidth-intensive video-based and distributed computing applications. Distributed supercomputing applications would require even greater increases in network bandwidth. Second, the networks will have to be more flexible than today's networks—they will be supporting a more diverse range of services, with a wider range of bandwidth requirements.

HIGHER BANDWIDTH

The bandwidth requirement for each type of application depends on a number of factors.

[7] Office of Science Technology Policy, "Grand Challenges 1993: High Performance Computing and Communications," p. 15.

[8] Catlett, op. cit., footnote 6, pp. 46-49.

Box 3-B—Distributed Supercomputing

Supercomputer-based simulations are becoming essential tools for science and engineering. Often, scientists are able to study problems that would be difficult or impossible to study theoretically or experimentally. For example, a number of researchers are developing climate models that can be used to predict the evolution of the Earth's climate over the coming decades. Computational science is becoming more widely used as the increasing power of new supercomputers brings more problems within reach.

One of the goals of the HPCC program is the development of computer technology that will allow scientists to tackle problems that are beyond the capabilities of today's machines. Some problems simply take too long to compute—some current models of the ozone depletion process take 10 hours of supercomputer time to compute the complex chemical reactions that take place in every day of "real" time. Other problems cannot be studied at a useful level of detail—computer power might limit a climate model to tracking the evolution of temperatures at only a limited number of points on the globe.

Greater processing power is also required for "interactivity," completing a computation in time to allow a user to take some action to control an instrument, change the parameters of a simulation, or "browse" other data sets in a database (see box 3-A). This requires that the computation of the model and the visualization processing be done in a fraction of a second. Today, images can take a considerable amount of time to compute, and are generally processed "off line" after the computation has been completed. Because of the time required to process newly computed or collected data, much of the data often goes unused.

The testbed program is investigating the use of gigabit networks to help address difficult computational science problems. High-speed networks may enable increased processing power, by linking several computers through the network. For example, a model could be computed on a supercomputer and then sent to a special graphics processor for the visualization processing, or a model could be split into two parts, with two supercomputers working in parallel to solve the problem. Networks also allow data from multiple sources to be used in a computation—large databases and scientific instruments, for example.

In the testbeds described in chapter 4, distributed supercomputing is used to increase processing power to study long-term weather models (part of the CASA testbed research), molecular dynamics (NECTAR), and chemical modeling (CASA). The use of networked computers to speed up the visualization process in an interactive fashion is being explored as part of applications for medical treatment planning (VISTAnet) and radio astronomy (BLANCA). Navigation of multiple large databases and associated visualization are used for terrain visualization (CASA), atmospheric sciences (BLANCA), and terrain navigation (MAGIC).

The testbeds are also working on the systems software and "tools" that will support these applications. Today, implementing distributed applications requires detailed knowledge of the behavior of the network and the characteristics of different computers. Distributed supercomputing will only be widely used by scientists if they can be freed of the need to learn these details, and can concentrate on the science aspects of their simulations. The testbeds are developing software modules that implement commonly used functions, and programs that automate parts of the applications development process. In the long run, the objective is to create software support and distributed operating systems for a "metacomputer," which would hide the complexity of networked computers and appear to the programmer as a single computer.

SOURCES: Gary Stix, "Gigabit Connection," *Scientific American*, October 1990, p. 118; Matthew Arrott and Sara Latta, "Perspectives on Visualization," *IEEE Spectrum*, vol. 29, No. 9, September 1992, pp. 61-65; Larry Smarr and Charles E. Catlett, "Metacomputing," *Communications of the ACM*, vol. 35, No. 6, June 1992, pp. 45-52; Office of Science and Technology Policy, "Grand Challenges 1993: High Performance Computing and Communications," p. 54.

Because of advances in "compression" technology, it now appears that relatively modest increases in bandwidth can accommodate many simple video and multimedia applications. There are many ways to convert a video signal to a digital stream of bits; new compression algorithms are able to squeeze the information content into fewer bits without significantly affecting picture quality.[9] These improvements have resulted from a better understanding of the mathematics of signal processing and also from research on how people perceive images.[10] In addition, increased processing power due to advances in microelectronics has allowed sending and receiving computers to do more complex signal processing.

Advances in compression technology have been dramatic. While it was once believed that a 155 Mb/s fiber optic link could carry only a single high-definition television (HDTV) signal, it is now believed that such a link can carry five or six HDTV signals.[11] In addition, it now appears that many simple video and multimedia applications will not require broadband fiber access to the network. New compression techniques are able to compress VCR-quality video to a few megabits per second, bandwidths that can be supported by new schemes for converting the telephone companies' existing copper local loops to digital service.[12] There are a number of emerging video and image compression standards—the most prominent of these will be the HDTV standard to be chosen by the Federal Communications Commission in 1993.[13]

However, there are still many possible applications that would more fully use the capacity of fiber.[14] These are the kinds of applications that are being investigated by the testbeds described in chapter 4. One possibility is distributed supercomputing—the use of high-bandwidth links to combine the processing power of multiple supercomputers. There are also applications that require images or video of a quality that can be only supported by fiber, despite advances in compression technology. In some cases, such as some medical applications, compression cannot be used because it destroys vital data. Other applications may demand very high bandwidths because many medium bandwidth streams of data are delivered to the user at once, allowing the user to select, combine, or process the streams at the workstation.[15] "Telepresence" or "virtual reality" applications require the delivery of large amounts of data in order to create the illusion of a user being in a distant location.

FLEXIBILITY

The second requirement the envisioned applications place on advanced networks is flexibility. First, new network technologies should be sufficiently flexible to carry all kinds of traffic. The integrated services concept envisions a network in which the same links and switches are used, to the extent possible, for all types of traffic. Integrated services networks may be more efficient than separate networks, and would also match advances in computer technology that allow computers to run multimedia applications. Today, different network technologies are used for voice, video, and ordinary data traffic. As new services were required, new types of networks were constructed. The telephone network was

---

[9] Mark Robichaux, "Need More TV? TCI May Offer 500 Channels," *Wall Street Journal*, Dec. 3, 1992, p. B1.

[10] P. H. Ang et al., "Video Compression Makes Big Gains," *IEEE Spectrum*, October 1991, pp. 16-19.

[11] See, for a table of bandwidth requirements for compressed signals, J. Bryan Lyles and Daniel C. Swinehart, "The Emerging Gigabit Environment and the Role of Local ATM," *IEEE Communications*, vol. 30, No. 4, p. 54.

[12] Charles F. Mason, "Bell Atlantic Stretches Copper for Video Trial," *Telephony*, Oct. 26, 1992, p. 10.

[13] Mark Lewyn, "Sweating Out the HDTV Contest," *Business Week*, No. 3306, Feb. 22, 1993, pp. 92-94.

[14] Ransom and Spears, op. cit., footnote 5, pp. 30-40.

[15] Lyles and Swinehart, op. cit., footnote 11, p. 55.

augmented first by packet switched networks for data and then by a variety of specialized networks for video communication[16] and distribution, such as cable television networks. Separate networks were required in part because no switching technologies worked equally well with all services.

Broadband networks should also be able to accommodate a range of applications bandwidths, from the very small amount of bandwidth required for ordinary electronic mail to the gigabit rates needed for distributed supercomputing. Some kinds of switching technologies are more flexible than others in accommodating different bandwidths in the same network. Circuit switches, the type used in the current telephone network, limit applications to a small number of predetermined bandwidths, while packet switches are more flexible.

Flexibility is also important from a network planning standpoint. While there are some general ideas about the ways in which broadband networks will be used, there is no real operational experience. Ideally, the network technology that is deployed would be able to accommodate a range of different scenarios, and its effectiveness would not depend on network planners knowing the exact mix of future applications in advance. In addition, future networks will have to support a more diverse range of users, each with different bandwidth and service requirements. Network operators would like to deploy network technology that could provide services to a diverse range of users with a minimum amount of customization.

FAST PACKET NETWORKS

A number of new concepts for network design may meet the requirements for flexible broadband integrated services networks.[17] There is general agreement that these networks will rely on fiber optic transmission, which has sufficient bandwidth to carry video and other types of bandwidth-intensive services. There is also general agreement that future networks will use a concept called "fast packet switching," which provides the necessary processing power to keep up with increases in link bandwidth and the necessary flexibility to support different kinds of services and a range of bandwidth requirements.

Fast packet networks overcome the main weakness of traditional packet networks by using special control mechanisms to provide the consistent network performance required for video and other real-time services. In traditional packet networks such as the current Internet, the network could become heavily loaded in a way that degraded these services. Researchers are looking at a number of different schemes to either prevent networks from becoming too heavily loaded, or to minimize the effects of a heavily loaded network on traffic such as video that is sensitive to network performance. Fast packet switches can then act as the foundation for integrated services networks.

Both the computer and telecommunications industries are investigating fast-packet approaches. In response to the emerging consensus for these technologies, considerable work has been done on the development of the necessary network components such as switches. However, until recently most experience with these networks had been confined to relatively small-scale experimentation with local area networks, or simulation and mathematical modeling. One of the main purposes of the testbeds described in chapter 4 is to demonstrate these networks in a realistic environment. In addition to the testbeds, a number of other experimental fast packet networks are now being planned or are operational.

Two different kinds of fast packet switching are being studied by the testbeds. The most

[16] For example, the CONCERT network in North Carolina uses high-bandwidth microwave links to support videoconferencing and "teleclasses."

[17] Nim K. Cheung, "The Infrastructure for Gigabit Computer Networks," *IEEE Communications*, vol. 30, No. 4, April 1992, p. 60.

prominent of the fast packet switching concepts was first championed by the telephone companies and is called Asynchronous Transfer Mode, or ATM.[18] ATM has been chosen by the telecommunications industry's international standards group, the CCITT, as the foundation for the Broadband Integrated Services Digital Network (B-ISDN), a blueprint for the future development of the telephone network.[19] B-ISDN envisions the provision of 155 Mb/s or 622 Mb/s fiber optic access links to each customer, which would then be used to carry voice, video, and data traffic to support a range of applications.[20]

One of the most significant aspects of ATM is that it has subsequently been adopted by many companies in the computer industry, and by manufacturers of equipment for local area networks and private networks. This convergence with telecommunications industry plans[21] may simplify the task of connecting different kinds of networks—in the past, local and wide area networks have used different technologies. However, technologies other than ATM have also been proposed for local area networks. Most of the testbeds described in the next chapter are using supercomputer industry networking standards that require the construction of modules that convert between the supercomputer network format and ATM.

While most packet networks use packets that can be very long and vary in size depending on the data being carried, ATM networks use short packets called "cells" that are always the same length (figure 3-1). If an ATM network is being used to carry Internet traffic, the Internet packets would be broken into a series of cells (figure 3-2(a)). After traveling through the network, the cells would be reassembled into the Internet packet and delivered to the destination computer. The same network could also carry video or sound: as the video or sound was digitized, the computer would load the bits into a cell (figure 3-2(b)-(c)). As soon as the cell was filled, it would be sent into the network and the user would begin filling the next cell. The cells carrying the video and Internet packet data would travel through the network together, sharing the same links (figure 3-3) and being processed by the same switches.

The second approach to fast packet switching being studied by the testbeds is called Packet Transfer Mode or PTM. The version being studied in the testbeds has not been adopted by standards committees. PTM does not use short cells, but more traditional packets that can be longer if necessary. This may simplify the task of carrying long Internet packets, because the computer does not have to break up the packet into many cells. ATM may also encounter problems at very high bandwidths—because the cells are so

Figure 3-1—ATM Cell

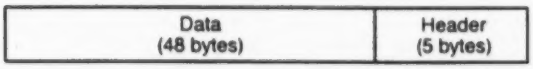

International standards specify that ATM cells are 53 "bytes" long, 48 bytes to carry the applications data and 5 bytes for the header, which is used to carry such information as the destination's address. By contrast, the packets used in traditional packet networks can be several thousand bytes long. A "byte" is the computer science term for eight bits.

SOURCE: Office of Technology Assessment, 1993.

[18] Robert W. Lucky, Executive Director, Communications Sciences Research Division, AT&T Bell Laboratories, testimony at hearings before the Joint Economic Committee, June 12, 1992.

[19] A. Day, "International Standardization of B-ISDN," *IEEE LTS*, vol. 2, No. 3, August 1991.

[20] Peter Bocke and Heinrich Armbruster, "Broadband Services: An Overview," *Telecommunications*, vol. 25, No. 12, December 1991, pp. 24-32.

[21] John J. Keller, "AT&T Sets Alliance to Make Gear to Provide Multimedia Services," *Wall Street Journal*, Jan. 13, 1993, p. B6; John McQuillan, "Who's Who in ATM," *Business Communications Review*, August 1992, p. 10.

Figure 3-2—ATM

(a) Data communications using ATM

Packets are broken into several ATM cells. After travelling through the network, the cells are reassembled into packets.

(b) Video communications using ATM

The digital video bits are put in cells and sent through the network. At the destination, the bits are removed from the cells.

(c) Voice communications using ATM

Voice is handled in the same way as video.

SOURCE: Office of Technology Assessment, 1993.

short, there is little time to process each cell before the next one arrives. However, ATM proponents believe that the use of cells makes it easier to develop the control mechanisms that support real-time traffic such as video, and to be better suited to voice traffic. One purpose of the testbed research is to compare the two approaches to fast packet switching with realistic traffic.

NETWORK COMPONENT DEVELOPMENT-CURRENT STATUS

The telecommunications and computer industries have been working intensively to develop components for fast packet networks. The components are in varying stages of development. Fiber optic transmission links are the most advanced in their development—very high bandwidth optical

Figure 3-3—Integrated Services Using Fast Packet Switching

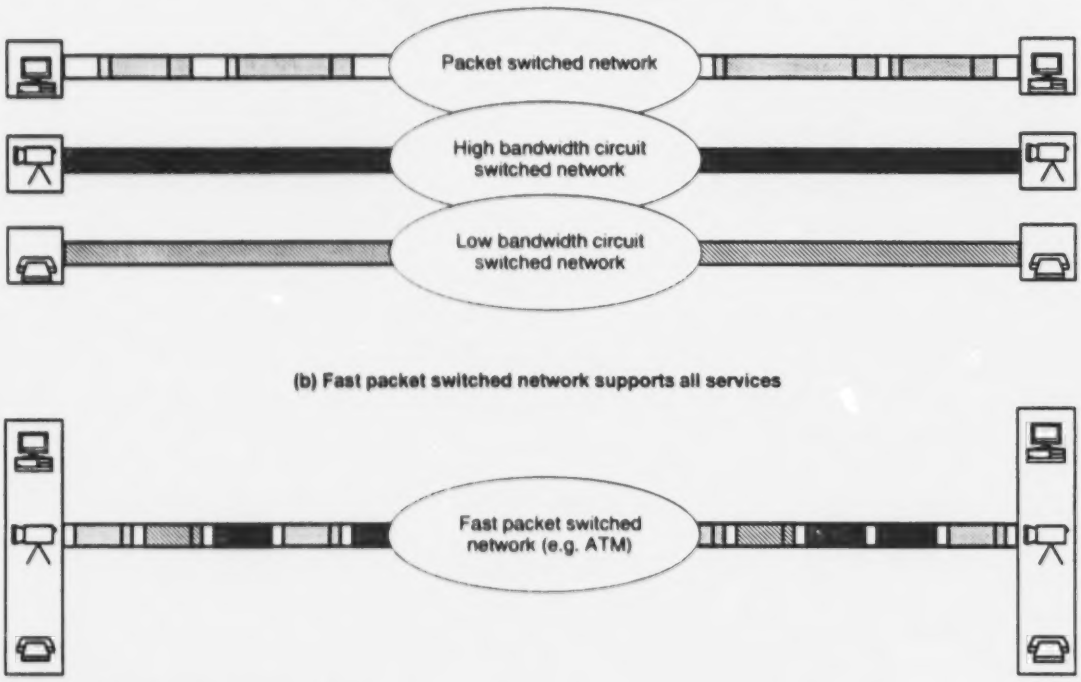
(a) Today, separate networks for each service
Packet switched network
High bandwidth circuit switched network
Low bandwidth circuit switched network
(b) Fast packet switched network supports all services
Fast packet switched network (e.g. ATM)

SOURCE: Office of Technology Assessment, 1993.

transmission systems are now commercially available. Fast packet switches are the subject of considerable industry research and development; most of the major telecommunications industry suppliers have had intensive ATM switch development efforts since 1987 or 1988, when it became clear that standards groups were going to adopt ATM. Some fast packet switches are becoming available commercially, but switches are less advanced in their development than the fiber optic links. Important work also remains to be done on the design of the software and hardware that handle the connection between the computer and the network.

Optical Fiber

Optical fiber has clearly emerged as an enabling technology for broadband networks. With increased bandwidth the links will be able to move data more quickly and support the transport of bandwidth-intensive traffic such as video. The development of the transmission equipment that handles gigabit rates is no longer a research issue. Although configured to support voice telephone calls, many fiber optic links in today's telephone network operate at more than one gigabit per second. Furthermore, the fiber cable is already widely deployed in much of the telephone network, especially in the interoffice portions of the network that will provide most of the transmission facilities for the agency backbones and regional networks.

For the telephone company fiber links to support the gigabit NREN and other broadband services, new transmission equipment will have to be deployed. This equipment is expected to

conform to a new standard called Synchronous Optical Network (SONET), and is now becoming commercially available. While fiber has been used in the telephone network for a number of years, the link capacity was mainly configured to carry thousands of low bandwidth telephone calls. SONET transmission links, on the other hand, can be configured to support the high-bandwidth channels required for advanced networks. For example, the transmission facilities to be used in the testbeds employ a 2.4 Gb/s SONET link, which can be divided into four 600 Mb/s channels.

Switches

The development of fast packet switches is less advanced than the development of transmission links. However, there has been considerable theoretical work done on switch design, prototypes have been developed, and some early commercial products are becoming available. By the end of 1993 or early 1994, several suppliers should have ATM products on the market. The early products are designed primarily for private networks or carrier networks operating at 45 Mb/s or 155 Mb/s, not gigabit rates. 155 Mb/s is the bandwidth specified by the telecommunications industry's standards group, the CCITT, for the Broadband Integrated Services Digital Network (B-ISDN) service.

There are many different ideas for how to build fast packet switches—the "best" design depends on assumptions about the number of users, the bandwidth of the network, and the mix of traffic. However, all of the proposed switch designs rely on hardware, in order to speed the processing, and are usually based on "parallel" designs that allow many packets or cells to be processed at once.

If ATM switches do become central to telephone company networks, then there will be demand for large switches to replace the current "central offices" that handle tens of thousands of lines. Most of the ATM switches now becoming available only handle a small number of lines—16 or 64 lines are common configurations. Initially, ATM switches will probably be introduced to support new services, rather than as a replacement for existing central office switches.[22] Building an ATM switch that can serve thousands of lines is a difficult task, requiring further research on switch design and device technology and packaging. The move to ATM switching has the potential to change the market positions of telecommunications equipment manufacturers, much as the transition from analog to digital created market opportunities in the late 1970s.[23]

Switches control the flow of packets using considerable software "intelligence." For example, if the network is heavily loaded, a switch may decide to handle video or other performance-sensitive traffic first. Switches may also help prevent the network from becoming too busy—they may prevent a user from sending traffic, or verify that users are not using more than their share of the network capacity. These aspects of the control of the network are still important research issues, however; there are many different proposals for managing fast packet networks. As a result, some of the prototype switches used in the testbeds described in chapter 4 are flexible enough to allow researchers to reprogram the network control mechanisms.

Computers

The use of high-bandwidth links and switches will expose new bottlenecks inside many computers.[24] New computer designs may have to be

[22] Richard Karpinski, "Ameritech Issues End-to-end ATM RFP," *Telephony*, vol. 224, No. 9, Mar. 1, 1993, p. 14.

[23] John J. Keller, "Telephone Switching Moves Toward Increased Speed," *Wall Street Journal*, Nov. 4, 1992, p. B4; Steven Titch, "Northern Keeps ATM Details Slim," *Telephony*, Sept. 28, 1992, p. 13; "NEC Lands ATM Pact for Wiltel," *Telephony*, Oct. 26, 1992, p. 8.

[24] H.T. Kung, "Gigabit Local Area Networks: A Systems Perspective," *IEEE Communications*, vol. 30, No. 4, April 1992, pp. 86-88.

developed to improve the rate at which data can be transferred from the network, through the computer's internal circuitry, and into memory, where it can be used by the applications software. Both the internal circuitry and the memory of today's computers are limited in the rate at which they can transfer data. In the past, the design and operation of computers has focused on the task of maximizing the processing power once the data is in memory, not the larger problem of maximizing the performance of networked applications.

Computers may also require additional processors or hardware to process protocols. With low speed networks, the computer's main processor was powerful enough to handle the communications functions and still have enough time left to run applications. This may continue to be the case as the processing power of computers continues to increase. However, in some cases it may be necessary to relieve the processor of some of the burden of handling the communications functions. This is likely to be the case with ATM-based networks—because the cells are so short, there are many cells to be processed in a given amount of time. Special "network interfaces" that speed the protocol processing are being developed for a number of different computers as part of the testbed project discussed in the next chapter.

APPLICATION OF BROADBAND TECHNOLOGIES

The broadband technologies discussed in this chapter will be used in both the Internet and in other networks, such as private networks or the public switched network. The use of broadband technologies in the Internet is linked to their deployment in the public switched network in two respects. First, the Internet will probably continue to rely on the public network's transmission infrastructure. As a result, it is dependent on the rate of deployment and the cost of SONET links. Second, the carriers may use their new ATM-based infrastructure as a way to play a more active role in the computer communications business and offer Internet services.

Application to the Internet

The links in the high-speed networks in the core of the Internet are expected to use the SONET-based transmission infrastructure that the telephone companies are planning to deploy. SONET is actually a family of transmission rates—there is 155 Mb/s, 622 Mb/s, and 2.4 Gb/s SONET equipment becoming available now. Users that need access to the Internet at broadband rates will also use SONET for their access links. Large universities and commercial users of the Internet would be able to make special arrangements with their local exchange carrier for the provision of fiber access. The rate at which the carriers will more broadly deploy optical fiber in the local loop depends on the resolution of complex economic and policy issues.

However, many users of the Internet will not require fiber optic access links in the near term. The carriers have proposed several new technologies that would convert existing copper local loops to digital service. These technologies do not support true broadband capabilities, but still represent a significant improvement over the existing analog local loop. They include the Integrated Services Digital Network (ISDN) (see box 3-C for a description of ISDN), High-bit-rate Digital Subscriber Line (HDSL)[25], and Asymmetric Digital Subscriber Line (ADSL) standards. ISDN provides access at 144 kb/s; HDSL provides access at 768 kb/s. These technologies are available on a limited basis from the carriers and are the subject of a number of trials and demonstrations.[26] The pace of their deployment depends

[25] Gerald A. Greenen and William R. Murphy, "HDSL: Increasing the Utility of Copper-Based Transmission Networks," *Telecommunications*, August 1992, p. 55.

[26] Mason, op. cit., footnote 12.

Box 3-C—ISDN

ISDN (Integrated Services Digital Network) is a telecommunications industry standard for upgrading local loops to digital service. This "last mile" of the network, the wire that connects a telephone network to its customers, is less sophisticated than other parts of the network. The core of the telephone network uses high-capacity, digital, fiber optic links. The local loop, by contrast, uses low-capacity, analog, copper wires. This technology is acceptable for ordinary telephone service, but more sophisticated services will require upgraded local loops.

When work began on the ISDN standards in the mid-1970s, it was believed that ISDN would soon be deployed to all of the telephone network's customers. Today, ISDN is used only on a very limited basis, due to delays in completing the standards and several regulatory and economic questions. Because of the delays in deploying ISDN, large business customers found more capable technologies. More importantly, a new vision of the future of the local loop emerged—the wiring of homes with fiber optic links—and ISDN was no longer viewed as a technology with an important role to play.

There is now renewed interest in ISDN, however, as an "intermediate" step between the current analog local loop and the use of fiber optics. Because of the cost of deploying fiber, it may be many years before significant numbers of homes are connected. ISDN is cheaper than fiber, can be deployed sooner, and, while its capacity is only a fraction of fiber's, represents a significant improvement over the current analog local loop. While ISDN will not become the universal network standard once envisioned, it may play a role in providing better network access to certain groups of users.

For example, one possible application might be telecommuting, which allows employees to work at a desktop computer at home. To connect to the office computers, workers today would need a device called a modem, which lets them send digital computer data over the analog local loop. Common modem standards transmit data at 2,400 or 9,600 bits per second; ISDN, by contrast, provides two 64,000 bits per second (64 kb/s) channels. This would allow videoconferencing of reasonable quality, faster transfer of graphics information, and better quality fax transmission. It would also permit much-improved access to the Internet for home users. Today, good-quality access to the Internet is usually only available to large customers who are able to arrange for special digital access lines to be provided by their local telephone company.

ISDN allows the existing copper local loops to be used for digital service. However, it requires users to buy new equipment for their end of the line, which converts their data to the ISDN format. It also requires that the telephone company's equipment, such as the "central office" switches, be upgraded. Currently, the user equipment is expensive and only one-third of the telephone lines are connected to switches that are "ISDN ready." In addition, ISDN communications are hampered by the fact that different equipment manufacturers have implemented their own versions of ISDN, despite the fact that it was developed to be a standard. In most of the United States, ISDN is not available as a regular service.

However, some progress is being made toward overcoming ISDN's problems. The industry has a number of initiatives that are intended to encourage the development of ISDN equipment that conforms to a common specification. The Regional Bell Operating Companies, which provide local telephone service in most of the United States, have announced that they are planning to make 56 percent of their lines ISDN-ready by the end of 1994. In addition, the cost of users' ISDN equipment may decline as it becomes more widely used.

Broadband ISDN, which is discussed on p. 46, uses very different technology from "ordinary" or "narrowband" ISDN. Narrowband ISDN is best viewed as a digital upgrade of the telephone network's copper local loop. Broadband ISDN, by contrast, requires fiber optics and Asynchronous Transfer Mode (ATM), a new approach to network design discussed in detail in this background paper. ISDN and Broadband ISDN have little in common other than their names.

SOURCES: James N. Budwey, "It's Time to Get Off the 'POTS'," *Telecommunications*, August 1992, p. 4; Cindy Skrzycki, " 'Data Highway' Plan Costs May Decline," *The Washington Post*, Jan. 20, 1993, p. G1; Bob Wallace, "Study Raises Concerns About National ISDN," *Network World*, June 29, 1992, p. 27; Mitchell Kapor, President, Electronic Frontier Foundation, testimony at hearings before the House Subcommittee on Telecommunications and Finance, Jan. 19, 1993; Steve Lohr, "Computer Makers Told to Get Involved in Rules," *The New York Times*, Feb. 23, 1993, p. D2.

on resolving standards issues and on business decisions made by the carriers.[27]

Internet traffic may be handled by some of the new fast packet switching systems. As was noted above (p. 43), fast packet networks can carry Internet traffic if the Internet packets are first converted to the appropriate fast packet format—for example, if the Internet packets are broken up into a series of ATM cells. It is likely that other types of switching technologies will also be used. The success of the Internet is due in substantial part to the commonality of protocols that support the technological diversity of the interconnected networks. Some networks will continue to use "routers," similar to those used in today's backbone networks, while others may employ the new fast packet switching technologies or some of the new data communications services that the carriers may offer, such as Frame Relay or SMDS.

Public Network

In many ways, the most significant aspect of ATM is that it was first championed by the telephone companies and is now a key component of telephone company planning. ATM represents a dramatic change in the design of telecommunications industry networks. Traditionally, the industry has not used packet switches. It used the circuit switches that were ideally suited to carrying voice telephone traffic. The industry standards group, the CCITT, chose ATM because it believed that simply upgrading the existing circuit switched network to higher bandwidths would not provide the necessary flexibility to support future services. ATM is a central component of carrier strategy; they hope to use ATM as the basis for a range of future services, including video, Internet services, and other data communications services such as Frame Relay or SMDS.[28]

ATM's flexibility offers the carriers an opportunity to enter a variety of markets and quickly offer new services with a common infrastructure. However, some believe that ATM's flexibility also means that it is a compromise technology, and that more specialized network technology will continue to play a role.[29] Moreover, there are still important economic considerations for the telephone companies as they determine the best way to evolve from the current network to an ATM-based infrastructure.[30] Both service providers and manufacturers are facing difficult decisions about the timing of their investments and the appropriate migration scenarios.[31] Deployment decisions depend on estimates of future revenues, equipment costs, the viability of competing technologies, and the carriers' investment in their existing networks.

---

[27] Edmund L. Andrews, "'Baby Bells' Rift Threatens An Advanced Phone Service," *The New York Times*, Dec. 1, 1992, p. D1.

[28] Ben Lisowski and Louise Reingold, "Sprint's Evolution to Broadband ISDN," *IEEE Communications*, August 1992, vol. 30, No. 8, pp. 28-30; John Williamson and Steven Titch, "Gazing Toward the Broadband Horizon," *Telephony*, Oct. 5, 1992, p. 38.

[29] Chin-Tau Lea, "What Should Be the Goal for ATM?" *IEEE Communications*, September 1992, vol. 6, No. 5, pp. 60-66.

[30] Stephen M. Walters, "A New Direction for Broadband ISDN," *IEEE Communications*, vol. 29, No. 9, September 1991, pp. 39-42.

[31] Carol Wilson, "It's Not That You Do It, But How You Do It That Counts," *Telephony*, June 15, 1992, p. 9.

Gigabit Research 4

As was discussed in chapter 3, the limitations of current networks and advances in computer technology led to new ideas for applications and broadband network design. This in turn led to hardware and software development for switches, computers, and other network components required for advanced networks. This chapter describes some of the research programs that are focusing on the next step—the development of test networks.[1] This task presents a difficult challenge, but it is hoped that the test networks will answer important research questions, provide experience with the construction of high-speed networks, and demonstrate their utility.

Several "testbeds" are being funded as part of the National Research and Education Network (NREN) initiative by the Advanced Research Projects Agency (ARPA) and the National Science Foundation (NSF). The testbed concept was first proposed to NSF in 1987 by the nonprofit Corporation for National Research Initiatives (CNRI). CNRI was then awarded a planning grant, and solicited proposals or "white papers" from prospective testbed participants. A subsequent proposal was then reviewed by NSF with a focus on funding levels, research objectives, and the composition of the testbeds. The project, cofunded by ARPA and NSF under a cooperative agreement with CNRI, began in 1990 and originally covered a 3-year research program. The program has now been extended by an additional fifteen months, through the end of 1994. CNRI is coordinat-

The HPCC program's six testbeds will demonstrate gigabit networking.

[1] Corporation for National Research Initiatives, "A Brief Description of the CNRI Gigabit Testbed Initiative," January 1992; Gary Stix, "Gigabit Connection," *Scientific American*, October 1990, pp. 118-119; John Markoff, "Computer Project Would Speed Data," *The New York Times*, June 8, 1990, p. A1; "Gigabit Network Testbeds," *IEEE Computer*, vol. 23, No. 9, September 1990, pp. 77-80.

ing five testbeds; a sixth testbed, funded by ARPA alone, was announced in June of 1992.

The testbeds are investigating gigabit networks, very high-speed broadband networks that represent the limit of what can be achieved today. Most current work on broadband networks is looking at lower bandwidths, such as the 155 Mb/s rate that will be used for the telephone companies' B-ISDN service. Because of the focus on gigabit rates, some aspects of the testbeds' research agenda are unique. In other respects, however, the testbeds are one of a number of research programs whose work will impact the NREN—fast packet switching technologies, for example, are being studied as part of many industry research projects.

RESEARCH OBJECTIVES

In general, the objective of the testbeds is to speed the deployment of advanced network technology, in the NREN and elsewhere. The networks are designed to provide a realistic test environment for the technologies outlined in the previous chapter. The switches and transmission equipment conform to emerging industry standards wherever possible. More speculative concepts such as optical switching are not being investigated by the testbeds—the focus is on the network technologies that are central to near-term industry planning. One purpose of the testbeds is to look at unresolved research questions. However, the most valuable aspect of the testbeds will be to demonstrate the feasibility of these networks and provide experience with their construction.

While much of the research is related to near-term industry plans, the testbeds are also looking into the future. The testbed networks achieve the highest bandwidths possible, given the constraints of emerging industry standards, current technology, and the time horizon of the program. The equipment used in the testbeds had to be such that it could reasonably be expected to be working in time to integrate the components and begin testing the networks by the end of the project. The applications are the most bandwidth-intensive possible, "gigabit applications" that require a full gigabit of bandwidth for each user. For the most part, these are distributed supercomputing applications that use the network to combine the processing power of multiple supercomputers.

The research is also related to the expected use of the network technology in the NREN environment. This emphasizes the use of Internet protocols with the new fast packet switching technologies, because the NREN program is linked to the evolution of the Internet. In addition, supercomputer-based applications of the type being investigated by the testbeds will play an important role in the gigabit NREN. However, not all issues relevant to the future development of the NREN are addressed by the testbeds: because of the emphasis on high-speed applications there is little work being done on applications that will be used outside the supercomputer community. Nor is there significant work being done on topics related to the growing size and complexity of the Internet (see ch. 2, p. 26, and ch. 5, p. 70).

Given the objective of demonstrating the feasibility of the emerging network design concepts, the testbeds are emphasizing the construction of working networks—much of the prior network research used modeling or simulation in "paper studies." Because there is little real experience with broadband networks, these models and simulations are based on assumed traffic patterns that may not be accurate. The testbeds are addressing this problem by building test networks and investigating both network and applications research simultaneously. The applications will provide a source of traffic with which to test the network components and protocols.

In addition, there is a focus on overall systems performance. The overall performance of a network depends on how well the individual components work together, not solely on the performance of any single component. In the past, researchers have tended to focus on the design of individual components; for example, some have

looked mainly at switch design, others at transmission systems, and others at protocol issues. In part, this has been due to the complexity of organizing research programs such as the testbeds that draw on the collaboration among several disciplines.

The five CNRI testbeds are AURORA, BLANCA, CASA, NECTAR, and VISTAnet, and are discussed in more detail in boxes 4-A to 4-E. The sixth testbed, MAGIC, is described in box 4-F.

Testbed Design

Each testbed is building a high-speed network that addresses wide area networking issues. The networks connect three or four sites—industry research laboratories, universities, Federal laboratories, and supercomputer centers—separated by anywhere from about 30 to many hundreds of miles. The focus on wide area networks provides a realistic testbed for the agency backbones and the public switched network. In the past, much of the research done on advanced networks has involved small "local area networks." These served to demonstrate the basic concepts and could be investigated by a small research group within a laboratory. The development of high-speed wide area networks is much more difficult, both technically and organizationally.

The testbed networks reflect the basic technology trends outlined in the previous chapter. The networks all use optical fiber transmission and fast packet switching. There is major emphasis on the use of the telephone companies' Asynchronous Transfer Mode (ATM) concept—five of the six testbeds use ATM in some fashion. One of the testbeds also uses Packet Transfer Mode (PTM), a second kind of fast packet switching, and is investigating the relationship between ATM and PTM. Industry standard equipment is used wherever possible—the transmission links conform to the current version of the Synchronous Optical Network (SONET) standard, and the switches and other components that process the ATM cells conform as closely as possible to the current versions of the international standards.

In order to focus on the systems issues, an effort was made to draw on component development work that was already underway when the testbed program started in 1990. This would limit the extent to which components had to be specially developed and allow more time to experiment with protocols, applications, and other issues related to the operation of the overall network. Because fiber optic technology is the most advanced part of the system, the testbeds are able to use early production models of SONET transmission equipment, operating at 622 Mb/s or 2.4 Gb/s. The switches, on the other hand, are mainly prototypes, as are the interfaces between the computers and the networks—before the testbed work focused attention on the issue of interconnecting different network elements, network interfaces received less attention than such areas as switch or protocol design.

At each testbed site are computers, switches, and network equipment. Computing resources available on the testbeds include workstations, vector supercomputers, massively parallel supercomputers, and some specialized processors. In some cases this equipment is connected directly to the wide area network; in other cases it is connected through a local area network. The local area networks are using newly emerging gigabit-per-second standards such as the supercomputer community's High Performance Parallel Interface (HIPPI) or pre-standard experimental technologies. A number of different interface devices are being developed to handle the conversion between the local area and wide area network protocols, especially the HIPPI to ATM conversion.

Of particular interest is the investigation of the use of networks to enable collaboration between scientists and bring to bear increased processing power on a scientific simulation. Many of the applications also use the network to support visualization or interactive control of a simulation executing on a distant computer. Scientists and

Box 4-A—AURORA

The AURORA network links four sites in the Northeast: the University of Pennsylvania in Philadelphia; Bell Communications Research (Bellcore) in Morristown, NJ; IBM's T.J. Watson Research Center, in Hawthorne NY; and the Massachusetts Institute of Technology (MIT), in Cambridge, MA (figure 4-A-1). Bellcore is the research arm of the Regional Bell Operating Companies (RBOCs) that provide local telephone service in much of the United States.

The testbed sites are connected by 622 Mb/s SONET channels. The transmission facilities are provided by three different carriers: interexchange links are provided by MCI, local exchange links to IBM and MIT are provided by NYNEX, and local exchange links to the University of Pennsylvania and Bellcore are provided by Bell Atlantic.

Figure 4-A-1—AURORA

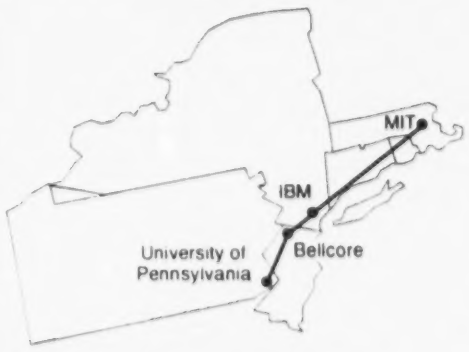

SOURCE: Office of Technology Assessment, 1993.

Each node will have experimental fast packet switches, which can either route traffic to a local area network on the testbed site or to another node. The local area networks will then distribute traffic to workstations, video monitors, and other devices. A number of network interfaces have been built to allow the workstations to connect to the local area networks and SONET transmission links. Bellcore and IBM are also supplying equipment for use in multimedia and videoconferencing applications.

AURORA is unique in two respects. First, it will employ two different switching technologies. Bellcore is contributing an ATM switch, based on the telecommunications industry standard that uses small, fixed length packets called cells. IBM is contributing a switch based on a second fast packet switching technology called Packet Transfer Mode (PTM) (part of IBM's "plaNET" network architecture). The PTM switch was designed to support a network architecture based on variable sized packets; it can, however, also handle ATM cells.

One of the research issues will be to compare the two types of switching technologies and to explore ways in which the two technologies can work together. In the current Internet, networks based on a wide variety of underlying technologies are used. Because both PTM and ATM may be used in future networks, it is important to gain understanding of how traffic could best be exchanged between these two networks. This work represents an initial step towards gigabit internetworking.

AURORA is also unique in that it is not investigating distributed supercomputing applications. Instead, it emphasizes experimentation with high-speed "multimedia" applications. Because video streams do not in general require a full gigabit of bandwidth, one concept is to deliver a gigabit stream consisting of a large number of medium-bandwidth video signals. For example, the network could be used to support an electronic classroom in which a user could select from different views of a classroom demonstration.

SOURCES: Biersack et al., "Gigabit Networking Research at Bellcore," *IEEE Network*, vol. 6, No. 2, March 1992, pp. 42-48; Cidon et al., "Bandwidth Management and Congestion Control in plaNET," *IEEE Communications*, vol. 29, No. 10, October 1991, pp. 54-64.

other researchers are developing applications in a number of areas, such as climate modeling, chemical modeling, and space science. Because, in the long run, scientists will want to develop applications without having to learn all of the details of the network and computers' operation, a number of modules and programs are being developed that simplify the task of applications development in a distributed computing environment.

The protocols generally conform to the existing Internet protocols, the protocols that will be the most widely used in the NREN. The use of well-understood, standard protocols also allows applications researchers to concentrate on applications development. The testbeds will provide a way to test the behavior of the Internet protocols in high-speed networks and to explore their use in a fast-packet-switched environment. However, the testbeds will also be testing a number of experimental protocols that may perform better with new network technologies. This research may serve to test ideas that will be incorporated in the Internet protocols in the future.

Testbed Organization

One of CNRI's key roles has been to assemble the testbed teams. The testbeds draw on researchers in industry, universities, supercomputer centers, and Federal laboratories. Some researchers within the groups have experience with traditional telecommunications issues, while others are more familiar with issues related to the Internet or supercomputer networking. The testbed research is necessarily multidisciplinary. In particular, each research group involves both network and applications researchers. The applications researchers have experience with supercomputers, visualization, graphics, and a variety of scientific disciplines. Network researchers draw on expertise with switches, transmission equipment, protocols, signal processing, and computer architecture.

While regular meetings are held between CNRI and program managers at ARPA and NSF, most of the responsibility for the management of the testbed program lies with CNRI. For example, one of CNRI's functions was to help develop the specifications for the transmission equipment that would be used in the testbeds. CNRI has also been responsible for maintaining the technical direction of the project, and has held a number of meetings on specific technologies. In addition, there have been annual meetings, which include attendees from a wider group than just the testbed participants, such as workstation manufacturers and government agencies, in an attempt to relate the testbed research to other industry activities and the broader NREN program.

One of CNRI's main contributions has been to ensure the participation of the carriers and other industrial partners. Participation of industry is essential to meeting the research goals of the project. First, the expertise required to develop many of the components required for high-speed network research is only available in industry. These components are complex, and their development involves the fabrication of custom integrated circuits and high-speed circuit design. Second, industry involvement has lowered the cost to the government of the program. The components developed by industry and the transmission capacity between the testbed sites have been contributed at no cost. Because of the contributions of industry, ARPA and NSF's support through the cooperative agreement with CNRI only covers a small part of the total cost of the project.[2]

There are a number of issues associated with the participation of industrial partners in the research venture. Some of these concerns are legal—there are antitrust issues, and further regulatory constraints govern the telecommunications industry. Another factor has been the competitive relationship among the testbed participants—while participating in the same research project, they are also competitors in various lines of business. For example, the wider use of more sophisticated telecommunications industry services may not necessarily be in the interests of companies that have emerged to offer computer networking services.

Moreover, some aspects of the research do not reflect industry priorities. Because of the cost of true gigabit access, it has been estimated that it

[2] Stix, op. cit., footnote 1, p. 118.

Box 4-B—BLANCA

The sites on the BLANCA network are more widely separated than those of the other testbeds. The network links AT&T Bell Laboratories in New Jersey, the University of Wisconsin and the University of Illinois, and the University of California-Berkeley and Lawrence Berkeley Laboratories (figure 4-B-1). Because of the cost of gigabit transmission facilities, high-speed links will initially be used only for some parts of the network. The cross-country segments of the network will use 45 Mb/s T3 links. While this bandwidth is not sufficient for distributed supercomputing applications, the BLANCA network will still provide an environment for researching the behavior of new protocols in a large network.

Figure 4-B-1—BLANCA

SOURCE: Office of Technology Assessment, 1993.

BLANCA is an ATM-based network. The experimental ATM switches and other hardware are being supplied by AT&T Bell Labs, the main industrial partner for BLANCA. BLANCA builds on preexisting research relationships between Bell Labs and the University of Wisconsin, University of Illinois, and UC-Berkeley. The switches are designed in such a way as to allow researchers at the universities to "take over" the network, to control the switches with computer programs that implement their experimental protocols. The network research interests are similar to those of others looking at ATM-based networks, such as congestion control and the behavior of Internet protocols in an ATM-based network, and is being carried out primarily at UC-Berkeley, the University of Illinois, and the University of Wisconsin.

BLANCA emphasizes distributed supercomputing applications, as do most of the other testbeds. The applications work is being done at the National Center for Supercomputing Applications (NCSA), the University

would not be generally available to commercial customers until about 2005.[3] Much of the research agenda focuses on higher bandwidths and more specialized applications than are expected to have near-term commercial significance for the telecommunications industry. Industry planning is oriented more towards medium-bandwidth multimedia applications—applications that require more bandwidth than can be supported by current networks, but significantly less than the gigabit/second rates required by the supercomputer community. For example, the telecommunications industry's ATM-based Broadband Integrated Services Digital Network (B-ISDN) standard envisions 155 Mb/s channels to each customer in the near term. Furthermore, many of the interesting issues related to the operation of fast packet networks can be studied with lower bandwidth networks, although a few issues may only become apparent at gigabit/second speeds.[4]

TESTBED PROGRESS

The major research results of the testbeds are still to come. Most of the networks are not

[3] M. Niel Ransom and Dan R. Spears, "Applications of Public Gigabit Networks," *IEEE Network*, vol. 6, No. 2, March 1992, p. 30.

[4] Leonard Kleinrock, "The Latency/Bandwidth Tradeoff in Gigabit Networks," *IEEE Communications*, vol. 30, No. 4, April 1992, pp. 36-40.

of Wisconsin, and Lawrence Berkeley Laboratories. A significant part of the work involves the development of software packages and modules that make it easier for scientists to use distributed supercomputing applications. For example, NCSA has been developing modules that handle many of the networking functions; these free scientists of the need to learn all the details of the network's operation—they can simply incorporate the modules in their applications. Another project is developing a digital library that allows the user to control the retrieval and processing of data—one of the programs that can be accessed by this digital library handles visualization processing, for example.

The applications under development as part of BLANCA could be viewed as prototypes for the Grand Challenge problems to be investigated under the HPCC program. One important aspect of these problems is that they will require collaboration between geographically dispersed researchers. The network and computing environment could support this collaboration by providing facilities for videoconferencing. On a more sophisticated level, researchers at NCSA have developed a program that permits collaborative investigation of data. It permits a researcher to highlight a feature in the data displayed on a workstation screen; researchers at other sites would then see the same feature highlighted on their displays.

The Grand Challenge problems will also involve very large data sets. Processing the data into image form is computationally intensive, especially when it is necessary to view the data interactively. The University of Wisconsin and NCSA are investigating the use of high-bandwidth connections from a scientist's workstation to a supercomputer to provide the necessary computational resources for visualization processing.

A radio astronomy application being studied as part of the BLANCA project is looking at issues involved in visualizing large data sets. Arrays of radiotelescopes collect the data, which is then sent through the network to a supercomputer. A user at a workstation connected through a high-bandwidth network to the supercomputer can control the processing of the raw data into images, which are then sent through the network to the workstation for display.

SOURCES: BLANCA Annual Report; Charles E. Catlett, "In Search of Gigabit Applications," *IEEE Communications*, vol. 30, No. 4, April 1992, pp. 42-51; Larry Smarr and Charles E. Catlett, "Metacomputing," *Communications of the ACM*, vol. 35, No. 6, June 1992, pp. 45-52; Carolyn Duffy Marsan, "Gigabit Network at Siggraph Proves Need, Reveals Limits," *Federal Computer Week*, vol. 6, No. 22, Aug. 3, 1992, p. 1.

expected to be operational until the third quarter of 1993. After the initial planning stage, the testbed work during 1990-92 was mainly devoted to completing hardware development for the switches and interfaces, theoretical and simulation work on protocols, and development of the applications software and tools. The next step will be to integrate these components into a working network; this will occur in stages over the next few months. As the networks become operational, researchers will be able to begin addressing the unresolved research questions.

Work on the testbeds has been proceeding more slowly than expected. It had been hoped that there would be about a year to experiment with functioning networks before the end of the original 3-year program. Because most of the networks were not yet operational, a 15-month extension was granted in order to allow time to look at network-level issues and test the networks with applications. The delay has been due to the late availability of the transmission equipment and problems with the fabrication of switches and other hardware components.

Component Development

During the first 2 years of the testbed project, the participants have been working mainly on the completion of the individual network components. The SONET transmission equipment has taken longer than expected to become available, but is currently being tested and, in some cases,

Box 4-C—CASA

The CASA network connects four sites—the San Diego Supercomputer Center, Los Alamos National Laboratory, and the Jet Propulsion Laboratory and Caltech in Pasadena (figure 4-C-1). The links between the testbed sites are provided by MCI, Pacific Bell, and US West. Also participating in the project is the UCLA Atmospheric Sciences Department.

The main emphasis of the CASA project is on applications development—using the network to combine the processing power of multiple supercomputers. The three main applications under investigation are modeling of climate change, modeling of chemical reactions, and interactive visualization of data describing the Earth's crust. These applications all require more processing power than is available from a single supercomputer. For example, the CASA climate change model is limited to simulations of a decade or less in the current computing environment. One of the research issues concerns the partitioning of a computation among multiple supercomputers. While in theory multiple computers can be combined in order to solve a problem more quickly, the best way to allocate parts of the computation to different computers depends on a number of factors. It may be necessary, for example, to arrange the computation in such a way as to hide the time it takes for data to travel between the computers—even when traveling at speeds close to the speed of light, data can take a significant amount of time to travel from one computer to another. Efficient implementations would arrange the computation so that the supercomputer would be able to proceed with other calculations while waiting for data to arrive.

Figure 4-C-1—CASA

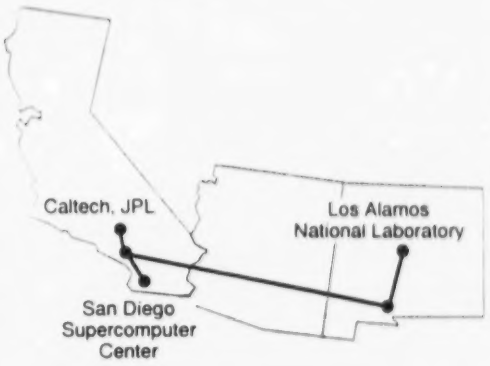

SOURCE: Office of Technology Assessment, 1993.

Efficient implementations may also be able to take advantage of the strengths and weaknesses of different supercomputer architectures. For example, researchers have determined that the climate modeling application can be split into a number of parts, each of which executes fastest on a particular kind of supercomputer. The part of the simulation that models oceans could be executed on a massively parallel computer, while the atmosphere would be modeled by a more conventional vector supercomputer. The two models would then exchange temperature information and other data at regular intervals. The CASA network provides access to a wide variety of supercomputer architectures, including different types of Cray Y-MPs, and massively parallel machines from Thinking Machines and Intel.

SOURCES: The CASA Gigabit Network Testbed, brochure; Charles E. Catlett, "In Search of Gigabit Applications," *IEEE Communications*, vol. 30, No. 4, April 1992, pp. 43-44; Elizabeth Smith, "NREN-Computer Expressway of the Future," *Gather/Scatter*, p. 3.

installed in the carrier networks. While the development of this equipment did not present any research issues, its availability was subject to factors affecting vendor development schedules. In part, these were hardware and software engineering issues. However, other factors have played a role; for example, the SONET equipment is very expensive and it is "high end" compared to the equipment that vendors expect to be the bulk of early demand. In addition, some aspects of the SONET standard have taken longer to complete than expected.

The development of the switch prototypes had been underway when the testbed work began, but in some cases the testbeds presented a more aggressive research target. The interfaces that

Box 4-D—NECTAR

The NECTAR network consists of a high-speed link that connects two local area networks, at Carnegie Mellon University and the Pittsburgh Supercomputer Center. The fiber links are being supplied by Bell Atlantic, and Bellcore and CMU are collaborating on the hardware design.

One area of research focuses on the interconnection of high-speed local and wide area networks. The NECTAR local area networks conform to a new standard called HIPPI (High Performance Parallel Interface), while the wide area connection between the two sites will use ATM cells over a SONET link. Research on this configuration is important because HIPPI is expected to be widely used by the supercomputer community, and the telephone companies are expected to deploy ATM- and SONET-based networks. A better understanding of the interactions between the two kinds of networks is expected to support future distributed supercomputing applications. As part of the research, a special interface circuit that converts between the HIPPI and ATM/SONET formats is being developed.

Another area of research is the implications of new high-speed networks for computer design. Most of today's computers were not designed in such a way as to optimize the task of moving data to and from the network. Applications may not be able to take advantage of improvements in the network bandwidth without improvements in the internal hardware or systems software of the computer. The NECTAR researchers are investigating different approaches to delivering data from the network to applications at high speed. Part of this effort has involved the construction of special "interface" circuits that free the computer's main processor of some of the protocol processing tasks.

Software development for the applications has been proceeding in parallel with the development of the hardware components. The applications have been tested in the local environment, and it is hoped that the applications can be made to work on the gigabit network with a minimum of modifications when it becomes operational. The applications are distributed supercomputing applications that take advantage of the combined power of multiple supercomputers to reduce the time needed to solve a problem. The NECTAR network will connect a number of different computers, including the workstations, the experimental iWarp parallel computer, and a variety of machines at PSC.

SOURCES: NECTAR annual report; H.T. Kung, "Gigabit Local Area Networks: A Systems Perspective," *IEEE Communications*, vol. 30, No. 4, April 1992, p. 79.

connect the computers to the network, or connect local and wide area networks, were designed specifically for the testbeds. Delays in the development of these components are due to their complexity and the demands of high-speed electronic design. A switch, for example, consists of a number of subsystems, each with a large number of standard and newly designed integrated circuits.[5] At the end of 1992, the custom integrated circuits had been designed, and most of the subsystems tested. The PTM switch to be used in the AURORA testbed has been completed, and the other switches and interfaces should be completed shortly.

To the extent possible, much of the work on protocols has been proceeding in parallel with the hardware development. This is expected to lead to faster research results once the networks become operational. Some of the work on protocols is conceptual and theoretical, and is done by simulation or by mathematically modeling the flow of data through a network. One of the main reasons for building the testbed networks is to test the assumptions that underlie these models and

[5] For a description of the components of one prototype switch, see Biersack et. al., "Gigabit Networking Research at Bellcore," *IEEE Communications*, vol. 6, No. 2, March 1992, p. 47.

Box 4-E—VISTAnet

One of the VISTAnet objectives was to use emerging public network technology and standards wherever possible. The switches and transmission equipment were supplied by the local telephone companies, Bell South and GTE, and are early production models from major equipment vendors. In contrast to the other testbeds, the switches are located not at one of the research sites but in telephone company central offices.

The VISTAnet network research involves Bell South, GTE, MCNC, North Carolina State University, and the computer science department at UNC-Chapel Hill. As with the other testbeds, a number of interface circuits have been developed. One of these boards also has the capability to collect data on traffic patterns, which will be used to develop more accurate traffic models for network research. This technology has been licensed, and may soon be available as a commercial product.

VISTAnet differs from the other testbeds in its focus on a single application. MCNC and UNC are collaborating on the development of an application that uses a network of powerful computers to help doctors plan cancer treatments. In radiation therapy, a number of treatment beams are used to deliver radiation to a tumor. Planning the orientation and strength of these beams is essential, because of the need to deliver radiation precisely and avoid the surrounding tissue. Planning involves calculating the distribution of radiation patterns for given beam strengths and orientations.

One limitation of existing technology is that the treatment planning is typically done only in two dimensions--doctors are only able to look at the distribution of radiation on a "slice" of the patient's body. The VISTAnet system, on the other hand, would be able to display data in three dimensions, giving doctors a much better view of the distribution of radiation.

Another limitation of current technology is that the planning process is not "interactive." Using today's technologies, interactivity is possible only at low resolutions—generating a high-resolution image takes too long and is done "off-line." VISTAnet is developing a system in which doctors can immediately see the effects of varying a parameter such as beam strength, allowing doctors to examine a greater range of treatment plans.

The interactive display of 3-D images of radiation distribution is computationally intensive. VISTAnet is using a high-bandwidth network to combine the processing power of a supercomputer, medical workstations, and a PixelPlanes machine, a special graphics processor developed at the University of North Carolina. A user at the workstation sends a description of the proposed beam strengths and orientation to a Cray supercomputer at MCNC, which then calculates the corresponding distribution of radiation. Data resulting from this computation is then sent to the PixelPlanes graphics processor at UNC-CH, which generates the image data that shows the radiation distribution as a 3-D image superimposed on the patient's body. The image data is then sent to the workstation for display. Much of the software has been developed and tested on low-speed versions of the VISTAnet network.

The VISTAnet application is a good example of the many different disciplines required to develop a distributed computing application. The medical component draws on expertise at the School of Medicine at UNC-CH. The application also requires the development of a user interface that allows doctors to rotate the image or highlight certain features. The graphics algorithms required to interactively generate 3-D volumes are themselves an important research area.

SOURCES: Dan Stevenson, ed., "VISTAnet Annual Report," April 1992; B.E. Basch et al., "VISTAnet: A BISDN Field Trial," *IEEE LTS*, vol. 2, No. 3, August 1991, pp. 22-30; M. Niel Ransom and Dan R. Spears, "Applications of Public Gigabit Networks," *IEEE Network*, vol. 6, No. 2, March 1992, p. 32; Daniel S. Stevenson and Julian G. Rosenman, "VISTAnet Gigabit Testbed," *IEEE Journal on Selected Areas in Telecommunications*, vol. 10, No. 9, December 1992, pp. 1423-1420.

Box 4-F—MAGIC

The MAGIC testbed is similar in many respects to the five CNRI testbeds, in that a high-speed network is used to provide access to supercomputing resources. As in the CNRI testbeds, there is considerable involvement of industry, the Internet protocols will be used, and the telecommunications services will conform to emerging industry standards like SONET and ATM. The application that will be used to test the network technology is of direct interest to the Department of Defense.

The participants in MAGIC are the Earth Resources Observation Systems Data Center, U.S. Army High-Performance Computing Research Center, the U.S. Army's Future Battle Laboratory, U.S. Geological Survey, Minnesota Supercomputer Center, SRI International, Lawrence Berkeley Laboratory, U.S. Department of Energy, MITRE, Digital Equipment Corp., the University of Kansas, Sprint, Southwestern Bell, Northern Telecom, and Split Rock Telecom.

The MAGIC network will connect four sites, the University of Kansas in Lawrence, Kansas, the U.S. Geological Survey in Sioux Falls, South Dakota, the U.S. Army's Future Battle Laboratory in Fort Leavenworth, Kansas, and the Minnesota Supercomputer Center in Minneapolis, Minnesota (figure 4-F-1). In the first phase of the project, the sites will be connected with point to point, 155 Mb/s or 622 Mb/s SONET circuits. In the second phase of the project, the network will use an ATM switch. The SONET and ATM services will be provided by Sprint.

Figure 4-F-1—MAGIC

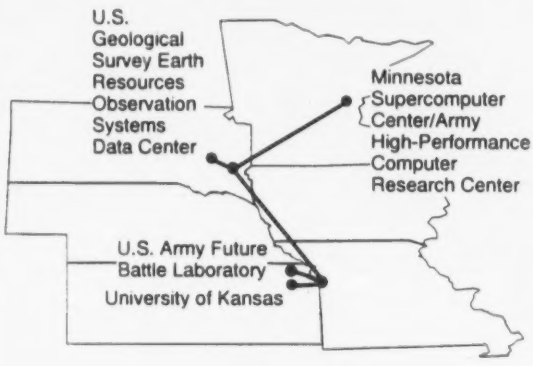

SOURCE: Office of Technology Assessment, 1993.

One of the research issues is the interconnection of different types of gigabit local area networks. Three different types of local area networks will be connected through the ATM wide area network. As part of the research effort, new modules will be built that convert from the local area network technology to ATM, and allow the interconnection of the different networks.

The application will investigate remote visualization of data drawn from a number of different sources. Information from a database at the U.S. Geological Survey will be sent through the network to a massively parallel supercomputer at the Minnesota Supercomputer Center. The supercomputer will compute images based on this data, and send the image data through the network to the Future Battle Laboratory, where it will be displayed on a workstation. The supercomputer provides the necessary processing power to select and view the images interactively (see the discussion of the VISTAnet application in box 4-E).

The test application will allow the simulation of walking or flying through a representation of a landscape. The Army believes that field officers could benefit from this capability, and that the application could be used for planning and educational activities. The landscape images are created from aerial images, satellite data, and geographic elevation data. Researchers will also study user interfaces to this type of application.

SOURCES: Anita Taff, "Sprint to Provide Services for DARPA Research Net," *Network World*, vol. 9, No. 27, July 6, 1992, p. 9; Tim Wilson, "Group Plots Gigabit Networking," *CommunicationsWeek International*, July 6, 1992, p. 35.

simulations. The protocol research also involves evaluating the behavior of existing networks like the Internet and writing software that will be used to program the switches, computers, and interfaces.

Work on the distributed supercomputing applications has also been proceeding in parallel with the hardware development. Much of the software development for the applications has been completed. In many cases, it has been possible to test these applications to a limited extent using existing high-speed local area networks or low-speed wide area networks like the Internet. Before writing the software, extensive analysis was done of the required computations, to determine how best to divide up the computations among the multiple computers that make up the overall system. Other important software development has involved the development of user interfaces and software tools that would make it easier to program distributed computing applications.

Systems Integration

The next objective of the testbed project will be to combine the network components into an operational network. This will begin once the transmission equipment is in place and work on the switches and other hardware has been completed. The systems integration task will proceed in stages, beginning with the simplest network possible, to minimize the number of sources of possible problems. VISTAnet began the integration process in the fall of 1992; the other testbeds should be in position to start this work by the third quarter of 1993. Over time, the networks will be expanded into more complex configurations.

The issues addressed in the early part of the systems integration phase are the low-level details of making sure that components designed by different groups work together or that a signal arrives in the format expected by a component's designer. These are the kinds of problems that are difficult to find when components are tested individually. For example, when the NSFNET backbone was upgraded from T1 to T3 links during 1990-92, the technical staff of the NSFNET backbone provider found that some components did not behave as expected under certain conditions, or unexpected traffic patterns required changes to the software and hardware.[6] Similar problems will probably be encountered as the testbeds begin to work through this stage with prototype or newly developed network components.

Network Research

One research issue concerns the algorithms used to control fast packet networks. These mechanisms are used to enable fast packet networks to support many different kinds of services using the same links and switches; one of the weaknesses of traditional packet networks was that they could not guarantee the kind of performance required for real-time applications such as video. In a fast packet network, software in the users' computers and in the switches will have to cooperate in managing the flow of traffic through the network in a way that supports all kinds of services. There have been many different mechanisms proposed for accomplishing this objective, but it is regarded as the most difficult problem with fast packet networks. The testbeds will provide an opportunity to test different control algorithms.

Another research issue is related to the development of distributed supercomputing applications. In these applications a computation is divided among multiple supercomputers; the network is then used to exchange data as the computation proceeds. Deciding how to allocate different parts of the computation to different supercomputers is a difficult problem. The best strategy depends in part on the characteristics of the network and the strengths and weaknesses of

[6] "T3 Upgrades to ANS Network Near Completion," *ANS Update*, vol. 1, No.1, 1991, p. 1.

different computers connected to the network—for example, some parts of a computation may be executed fastest on a massively parallel computer, while other parts may run faster on a vector computer. In order to maximize processing power, computers should not be idle while they are waiting for one of the other computers to finish its task or for data to be sent through the network.

Application of Testbed Research 5

The networking component of the High Performance Computing and Communications (HPCC) Program funds both research on gigabit technology and the deployment of this technology in the National Research and Education Network (NREN). One of the NREN's roles is to provide additional experience with advanced network technologies before they are deployed more widely in the national information infrastructure. However, the testbed research will also be applied directly to other networks, such as the common carriers' public switched network, without intermediate deployment in the NREN.

APPLICATION TO THE NREN

There is no overall NREN development plan; however, the National Science Foundation (NSF) is to coordinate the evolution of the Federal agency networks that are the core of the NREN.[1] During 1992, NSF, the Department of Energy (DOE), and the National Aeronautics and Space Administration (NASA) announced plans for the future development of their networks.[2] Based on these plans, the next-generation agency networks will likely be similar to the testbed networks, with an emphasis on Synchronous Optical Network (SONET) fiber optic transmission and fast packet switching. These broadband technologies are

The testbed research can be applied to networks other than the NREN.

[1] Office of Science and Technology Policy, "Grand Challenges 1993: High Performance Computing and Communications," p. 33.

[2] National Science Foundation, "Public Draft: Network Access Point Manager/ Routing Authority and Very High Speed Backbone Network Services provider for NSFNET and the NREN Program," June 12, 1992; James F. Leighton, Manager of Networking and Engineering, National Energy Research Supercomputer Center, Lawrence Livermore National Laboratory, "ESnet Fast-Packet Services Requirements Specification Document," Feb. 20, 1992.

being studied in a large number of research programs, but the testbed research is unique in its emphasis on building wide area gigabit networks and testing them with applications.

Agency Plans

The Federal agencies will not build their own "private" networks, but will obtain services from a network service provider. In effect, NSF, DOE, and NASA will act as early, large customers for advanced services. While industry has developed the switches and transmission equipment required for advanced network services, the agency backbones will be one of the first opportunities to integrate these components into a system that provides services to real users. Users of the agency backbones are knowledgeable about networking and will assist in integrating new network services with computers and applications to create useful systems.

The agency backbone services could be provided by a number of different organizations—the carriers, computer companies, or some of the emerging providers of commercial Internet services—or consortia. Provision of services for agency backbone networks provides valuable experience that the network operator may be able to translate into earlier availability of advanced services on a commercial basis.[3] For prospective players, the decision to participate in the provision of services to the agency networks weighs the experience gained and long-term strategic considerations against the cost of providing the service, which is greater than the money available from the Federal agencies.

To help stimulate market interest, DOE and NASA had originally decided to combine their NREN-related programs. A single supplier would have provided network services to both agencies, connecting sites such as DOE's Los Alamos National Laboratory or NASA's Ames Research Center. However, the General Accounting Office (GAO) overturned DOE's choice of contractor in March, 1993 (see ch. 1, p. 7). The steps that the agencies will take in response to this decision were still unclear at the time of publication, but it is possible that DOE and NASA will now decide to proceed separately. The procurement process has been significantly delayed, and will likely not be completed before the fourth quarter of 1993. Before the GAO decision, NASA and DOE had planned to begin connecting sites to the new network in mid-1993.

NSF issued a draft solicitation for its next-generation network in mid-1992. NSF plans to publish a final version of the solicitation and award a cooperative agreement during 1993. The new network is scheduled to begin operation in mid-1994. NSF's plans for the evolution of its network have greater implication for the evolution of the NREN and the Internet than do those of DOE and NASA. The current NSFNET backbone carries much more traffic than the other agency backbones [4] and serves a broader range of users. However, many of those users will not be able to use the next-generation backbone.

The new NSF network's primary purpose will be to connect the NSF supercomputer centers, enabling advanced distributed supercomputing applications. By contrast, today's NSFNET backbone is a "general-purpose" network that carries all types of research and education traffic. NSF envisions that in the future routine research and education traffic will be handled by commercial providers, not by the agency-operated backbone. There are a number of emerging commercial providers, and the network technology required to support routine traffic is sufficiently stable and reliable. This strategy would also free capacity on the backbone for experimental applications.

The transition to the new environment resulting from the changed role of the NSFNET

[3] Ann H. Lindstrom, "Sprint Blasts Ahead With ATM Deployment," *Telephony*, vol. 223, No. 8, Aug. 24, 1992.

[4] Stephen S. Wolff, Director, Division of Networking and Communications Research and Infrastructure, National Science Foundation, testimony at hearings before the House Subcommittee on Science, Mar. 12, 1992, Serial No. 120, p. 155.

backbone will require careful management to ensure stability. NSF's plan will affect significantly the existing three-level hierarchy of the NSFNET. The regional networks were designed to provide connections to sites on the current backbone, which in turn provides inter-regional connectivity. Under NSF's new plan, the backbone will serve many fewer sites and will no longer play the same central role in research and education networking. The regional networks will have to make new arrangements for interconnections and will be operating in a more competitive environment.

Agency Backbone Technology

The collaborative nature of the testbeds makes it more likely that the network technologies developed by industry will be suitable for operation in the agency backbone environment. The testbeds are emphasizing the technologies' use with the Internet protocols used by the agency networks, and are studying the interaction between fast packet networks and supercomputer network standards and applications. In addition, they emphasize the gigabit bandwidths required to support the Grand Challenge applications that are a key component of the overall HPCC program. The involvement of the carriers in the research program may also lead the carriers toward a more active role in providing NREN services.

While the plans for the evolution of the agency backbones are consistent with the target established by the testbeds, the agency networks will initially operate at lower bandwidths than the testbeds. The agency backbones will incorporate more of the technology from the testbed research as they evolve over time to meet the goal of the gigabit NREN. However, some issues cannot be addressed by the testbeds, or may be discovered only as the agency networks are deployed. Many of these issues are related to the more complex topologies (greater number of sites), larger number of users, and more diverse sources of traffic that will be present on the production networks.

TRANSMISSION TECHNOLOGY

The agencies envision the use of SONET equipment similar to that used in the testbeds, and have indicated that they hope to use 155 Mb/s SONET equipment in 1994 and then upgrade over time to 622 Mb/s (the next transmission rate in the SONET family) by 1996, the High Performance Computing Act's target year for the use of gigabit links. The 622 Mb/s rate, less than a full gigabit per second, is sometimes referred to as a "government gigabit."[5]

The rate at which the agency backbones will evolve depends on the timely deployment of a high-bandwidth SONET transmission infrastructure by the carriers. While much of the carriers' existing network uses fiber, SONET transmission equipment is required in order to support computer networking above the current T3 rates—it allows the fiber to be configured to carry high bandwidth channels. However, this equipment is extremely costly at this time and the carriers' deployment schedules have been slipping from earlier estimates.

The testbed networks will have also provided experience with the connection of supercomputers to high speed networks. "High end" users will require fiber links connecting their sites to the NREN. Only fiber is able to carry the large amounts of data needed for supercomputer-based applications. The testbeds are one of the first large-scale deployments of SONET to end-users, and considerable work has been done on interface devices to connect supercomputers and high-speed local area networks to fast packet switched networks. However, widespread use of high-speed networks will depend in part on the degree to which computer companies design their workstations to be fully integrated into a high-speed network. Today, bottlenecks encountered in mov-

[5] Carpenter et al., "Where Are We With Gigabits?" *IEEE Network*, vol. 6, No. 2, March 1992, p. 10.

ing data from the network into the computer's memory, where it can be used by the applications software, can limit the performance of the overall system.

SWITCHES

The next-generation backbone networks will use fast packet switching technology similar to that used in the testbeds. Initially, the switches will not be as sophisticated, because of the lower link bandwidths. The network operator's choice of switching technology, from among those being investigated in the testbeds and elsewhere, depends in part on long-term strategic considerations. If a carrier were to provide services for an agency network, it would probably use Asynchronous Transfer Mode (ATM) switches. ATM has been chosen as the foundation for the future development of carrier networks, and the provision of services for the agency backbone would provide an opportunity to gain experience with its use. Other providers might also choose to use ATM switches, or strategic considerations may lead to the choice of an alternate switching technology.

The DOE Reqest for Proposals issued in early 1992 specified the use of fast packet "cell relay" technology. "Cell relay" is a term used to describe both ATM and Switched Multimegabit Data Service (SMDS), a data communications service developed by the telephone companies. In the summer of 1992, DOE and NASA selected a contractor that proposed to use ATM. This DOE/NASA program would have been the first large-scale deployment of ATM. One of the goals of DOE and NASA is to encourage the development of commercial services by evaluating and demonstrating emerging technologies such as ATM. The agencies' effectiveness in performing this function may be reduced by any further delays resulting from GAO's decision overturning their choice of contractor.

The National Science Foundation's draft solicitation describing the evolution of its backbone network did not specify a particular type of switch.[6] NSF will allow prospective bidders to propose their choice of switching technology. The most likely option that would be proposed would be an ATM-based approach. Another type of fast packet technology, such as the PTM approach developed by one of the participants in the Aurora testbed, might also be used. The approach of "overlaying" an Internet network on a network that uses fast packet technology is not unique to ATM. However, ATM has broad support from industry standards committees.

OTHER NREN NETWORKS

The regional networks and other commercial providers of Internet services may also carry NREN traffic. Operators of these networks are faced with the same technology choices as those for the backbone networks. However, because many of these networks will require lower bandwidths than the backbones, they may continue to use "router-based" networks or use new "pre-broadband" services being offered by the carriers and other service providers. Two examples of these pre-broadband services are Frame Relay and SMDS.[7] These are packet switching services that can also be used to carry Internet traffic (see ch. 2, p. 34). Because the Internet protocols are able to hide differences in network technology from the users of the network, the NREN's networks can be based on a variety of different technologies.

Campus networks and other networks based primarily on local area networks will also become more capable. Local area network research is not currently a focus of the testbeds, although the interconnection of local and wide area networks

[6] National Science Foundation, "Request for Public Comment: Solicitation Concept NSFNET Backbone Network Services," June 12, 1992.

[7] Terry Sweeney, "TCP/IP-SMDS Trial Completed," *CommunicationsWeek*, Aug. 17, 1992, p. 25.

is being studied. One of the most important trends in local area network design is that there is a growing amount of support for ATM-based local area networks and products for ATM local area networks are beginning to appear. Other kinds of high-bandwidth local area network standards are also being studied by standards committees.

Applications

Because of the emphasis on gigabit applications, the testbed applications research is primarily applicable to high-end users of the NREN. The testbeds have been one of a small number of research programs to address supercomputer-related networking issues. These applications are, in general, of little concern to industry and would receive less attention without the testbeds. The testbeds' gigabit applications research will have important impacts on the overall HPCC initiative. Distributed supercomputing may be an important tool for bringing more processing power to bear on the Grand Challenge problems. In addition, the Grand Challenge teams will be scattered about the country and could use networks to support collaboration. The sizes of the data sets used in Grand Challenge problems will be very large, requiring high-bandwidth networks to move them from place to place within a reasonable period of time.

High-speed network support of supercomputing is important to the missions of the NSF supercomputer centers and the Federal laboratories. Led by testbed participants, the NSF supercomputer centers have proposed a concept that would make use of the distributed supercomputing ideas investigated by the testbeds.[8] They envision a "metacenter"—the use of the new high-speed backbone to integrate the computational and intellectual resources of the supercomputer centers.[9] In effect, it would be possible for the four supercomputer centers to act as a single center, distributing a computation among several machines as the computation required.

High-end users of the agency backbones are only part of the user community addressed by the NREN program. Few users will have access to a full gigabit/second of bandwidth, and the supercomputer applications studied by the testbeds are by definition highly specialized. For most users the primary result of improving network capability will be better performance with existing applications and the wider use of video and image-based communications. Because these capabilities may have considerable significance in commercial applications, much work is being done on these types of applications by industry. Some types of applications development, however, may require added support. Legislation introduced in the 103rd Congress (S.4 and H.R. 1757) seeks to expand support for applications development in a variety of education, medicine, manufacturing, and library settings.

Internetworking

The NREN is closely linked to the evolution of the Internet protocols that enable the thousands of independently operated networks that make up the Internet to exchange traffic. The testbeds are providing an opportunity to investigate the use of the Internet protocols in fast packet switched networks. The collaborative nature of the testbed program may be encouraging the Internet community to influence the ATM standards process to better support Internet protocols. In addition, the testbeds are investigating the behavior of the Internet protocols at high speeds, and comparing them to some new concepts in protocol design. In the past few years, a number of protocols have been proposed that may perform better in high-speed networks and are better suited to the new fiber-based, fast packet switched networks. For example, today's Internet protocols are designed

[8] John Markoff, "A Crucial Linkup in the US Data Highway," *New York Times*, Sept. 30, 1992, p. D8.

[9] Carolyn Duffy Marson, "NSF Pursues Computing Without Walls," *Federal Computer Week*, vol. 6, No. 35, Nov. 30, 1992, p. 1.

Figure 5-1—Growth in NSFNET Networks

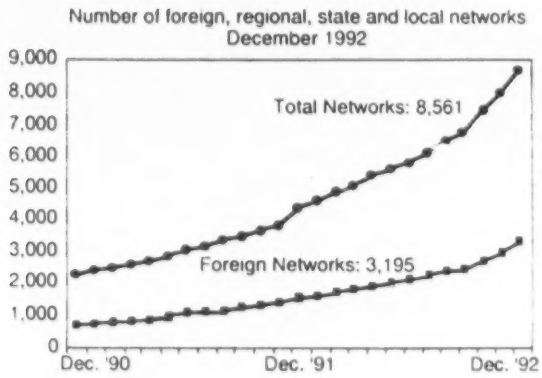

SOURCE: Merit, Inc., 1992.

to handle the types of transmission errors that occur with poor-quality copper lines, but rarely occur with new fiber-based transmission systems.

Other issues related to the evolution of the Internet protocols are not being studied by the testbeds. The main issue confronting the Internet community today is the growing size and complexity of the network—not increases in bandwidth. The growth in the number of users and networks that make up the Internet is putting pressure on current "routing" technology (figure 5-1). Routing is the process by which a path from one computer to another through a series of intermediate networks is determined. Calculating these paths using current algorithms demands a considerable amount of processing power; the problem is getting worse as the Internet continues to grow and become more complex. Routing issues have not been studied by the testbeds, which only connect a few sites.

Work on issues related to managing the growth of the Internet is being done primarily within the Internet community's technical organizations, such as the Internet Activities Board (IAB) and the Internet Engineering Task Force (IETF). The IETF consists of a number of working groups, one of which addresses routing issues. Currently, within the technical community there are many different proposals; some only address immediate problems, while others attempt to solve the problems in a way that will be satisfactory for a number of years. Besides addressing issues related to growth, some of the new routing algorithms may also take into account the growing diversity of service providers and network capabilities. Routing and management problems associated with the growing Internet are a major research area that requires more study.[10]

NSF's plan for the evolution of its network as part of the NREN program is linked closely to changes in routing technology. Today, the NSFNET backbone operator plays an especially important role in determining routes for research and education networks. As the Internet becomes more commercialized, however, it becomes less appropriate for NSF to be responsible for this aspect of its operation. NSF envisions reducing the reliance of Internet networks on the NSFNET backbone's operator for routing information.[11] NSF has proposed that the routing function be handled by a separate organization, the "routing authority," not by the operator of NSF's network. NSF's plan also calls for the creation of a number of Network Attachment Points (NAPs), where commercial networks and agency networks could obtain routing information and interconnect with each other (see box 5-A).

APPLICATION TO OTHER NETWORKS

The testbed program will also impact the evolution of the national information infrastructure more directly, without the intermediate stage of deployment in the NREN. This national information infrastructure includes the larger U.S. Internet—the NREN program targets only one part of the U.S. segment of the Internet (see

[10] ARPA is supporting research in these areas, such as through its DARTnet program.

[11] Robert Aiken et al., National Science Foundation, "NSF Implementation Plan for Interagency Interim NREN," May 1992.

figure 5-2).[12] It also includes a wide array of other services and technologies to be offered by the carriers, cable television companies, computer hardware and software companies, information service providers and others.

Application to the Internet

The Internet is increasingly expanding to serve communities other than the core research and education community that is the focus of the NREN program. The regional networks and new commercial providers now carry business traffic. The trends towards broader use of the Internet and growing numbers of users seem likely to continue. These will be driven in part by the advances in switching and transmission technology described in chapter 3. They depend to a greater extent on addressing the security concerns of commercial customers, the degree that use of Internet applications can be simplified, and the deployment of advanced digital local loop technologies.

It is possible that the switches and fiber optic links deployed by providers of agency backbone services will also carry commercial traffic. Some of the capacity would be used for the agency backbone network and some would be used to provide services to commercial customers. The Federal agency backbone would be the network's most important customer, acting as a catalyst for the deployment of the required switches and transmission equipment, while commercial customers would help to recover that portion of the costs of operating the network not covered by the Federal agencies' funding.

Other Services

The network technology studied by the testbeds is equally applicable to services other than Internet services. The research will also be applied directly to private networks, the common carriers' public switched network, and possibly cable television networks.[13] This is because the network technology used in the testbeds reflects near-term industry planning. While the testbeds have emphasized higher bandwidths and more specialized applications than have immediate commercial importance, the basic design of the testbed networks—such as the use of fast packet switching and SONET—reflects ideas that figure prominently in industry plans.

The carriers and other network operators could use the new advanced technologies to provide Internet services, or an array of other voice, video, and data communications services. Switch and transmission technologies, though advancing at different rates, are making substantial progress. Because of their commercial importance, fast packet and fiber optic technologies are being

Figure 5-2—Relationship Between NREN, Internet, and National Information Infrastructure

SOURCE: Office of Technology Assessment, 1993.

[12] For a discussion of the relationship of the Internet to the NREN, see chapter 2, p. 31.

[13] For a description of a cable provider's plans to use ATM, see Richard Karpinski, "Time Warner to LECs: Here We Come," *Telephony*, vol. 224, No. 5, Feb. 1, 1993, p. 8. IBM is testing its planET (PTM) fast packet switch in a cable network in Toronto, Canada. See "IBM's 1.2 Gigabit/sec. Networking Scheme," *Cable-Telco Report*, August 1992, p. 9.

Box 5-A—NSFNET Backbone Recompetition

The National Science Foundation's plans for the future development of its backbone network have attracted considerable scrutiny. The NSFNET backbone plays an especially important role in the Internet and in the National Research and Education Network program. Currently, NSF has a cooperative agreement with Merit Network, a not-for-profit organization of nine Michigan universities. However, Merit does not operate the NSFNET backbone "in-house." A second organization, Advanced Network & Services (ANS), operates the network—Merit obtains the services for the NSFNET backbone from ANS.

The cooperative agreement with Merit for the NSFNET backbone was announced in November 1987, and covered the 5-year period to November of 1992. Merit's proposal was submitted in partnership with IBM and MCI. The relationship between Merit and its partners changed in September of 1990, when Merit, IBM, and MCI announced the formation of ANS, as a not-for-profit corporation. ANS received capital from MCI and IBM at its formation, and IBM and MCI provide switches, transmission capacity, and other services to ANS. Over time, more of the responsibilities for the NSFNET backbone have been shifted from Merit to its subcontractor, ANS.

Over the life of the 5-year cooperative agreement with Merit, there have been three important changes in the Internet. First, Merit and ANS have increased the NSFNET backbone's link bandwidth from 56 kb/s, to 1.5 Mb/s (T1), to 45 Mb/s (T3). Second, the Internet has become a much more important part of the U.S. information infrastructure—the amount of traffic and the number of users has grown rapidly. Finally, the past 2 years have seen the emergence of commercial Internet service providers. In particular, ANS has created a for-profit subsidiary. While the T1 network was used only by NSFNET backbone traffic, the T3 network operated by ANS is shared by the NSFNET backbone and ANS's commercial customers.

The relationship between NSF, Merit, ANS, and other commercial providers was the subject of hearings before the House Subcommittee on Science in March of 1992. Concern was expressed by some witnesses that the current arrangement benefited ANS unduly, and had not been foreseen by the 1987 cooperative agreement with Merit. Other witnesses pointed to the success of the NSFNET backbone, the growth in the number of users, and the value of the equipment and services contributed by Merit and its partners.

Recompetition

In preparation for the expiration of the cooperative agreement with Merit in November of 1992, NSF studied a number of options for the future development of the NSFNET backbone. In studying these options, NSF had to take into account several factors that did not apply in 1987. One factor was the emergence of commercial providers. Any new plan for the backbone could not favor the incumbent, ANS, and would have to provide equal opportunity for all firms wishing to provide services to the NSFNET backbone. A second factor was the need for stability. The Internet is now an essential infrastructure for many more users than in 1987, and stability would have to be ensured during the transition to any new arrangement. Finally, NSF had to take into account the NSFNET backbone's central role in the NREN program.

One option studied by NSF was to discontinue direct funding of a backbone network. Instead, NSF could fund the regional networks and allow them to choose among commercial providers of interconnections, encouraging further development of the commercial networks. According to testimony of the director of the NSF division responsible for NSFNET, this plan was opposed by the regional networks and by other Federal agencies, in part because of concerns about stability during the transition to this environment.

As a result, NSF decided that it would continue to operate a backbone network. NSF's timetable called for extending the arrangement with ANS for up to 18 months beyond November 1992, to the middle of 1994. This eighteen month period was intended to allow time to 1) select the provider of the next-generation NSFNET backbone, and 2) install the required links and switches. Originally, NSF planned to make the awards in the middle of 1993, allowing a year for the transition to the new network.

The NSFNET Solicitation Concept

The Project Development Plan for the continued provision of NSFNET backbone services after the expiration of the agreement with Merit was published by NSF in November of 1991. This development plan stated the requirements for stability, fair competition, and support of NREN objectives. The Development Plan also presented the concept of splitting the current NSFNET backbone provider's tasks into two parts, and awarding each part to separate organizations.

NSF published a more detailed version of this plan in June of 1992 and requested public comments. According to the plan, one of the two awards would be for the provision of very high speed Backbone Network Services (vBNS). The vBNS provider would operate the links and switches and be responsible for moving packets through the NSFNET backbone. Among other requirements, the vBNS provider would establish a network that would operate at 155 Mbps or higher and would "provide for real-time multimedia services, including multicasting and video teleconferencing." NSF did not specify a switching or transmission technology; however, the reference to 155 Mbps implies the use of SONET transmission equipment.

The second award would be for the Routing Authority (RA). The routing authority would be responsible for the routing functions that had previously been performed by the backbone operator. The RA would also operate Network Access Points (NAPs), which would facilitate the connection of other networks to the vBNS and to each other. These could be other Federal networks, or commercial networks. The routing information required in order to facilitate the coordination of these networks would be stored in a database accessible at the NAPs. A total of about $10 million annually would be available for the two awards.

Changes to the Draft Solicitation

The public comments received by NSF in response to the draft proposal reflect the degree to which NSF's plans affect more than just the NSFNET backbone. NSF's proposed NAP/RA structure could best be characterized as an "architecture" for the NREN and the Internet, with significant implications for the larger information infrastructure. As such, the NSF's plans affected users, interexchange and local exchange carriers, regional networks and other current and prospective providers of Internet services, and other federally supported networks.

As of May, 1993, a revised version of the NSF solicitation had not been released. However, in December 1992, NSF outlined its intention to change its original plan in a number of ways. While the basic vBNS/NAP/RA structure was maintained, NSF indicated that it would make three awards, not two. The NAPs would not have to be operated by the Routing Authority, as had been specified in the draft solicitation, but could be operated by a separate organization.

More importantly, NSF announced that the new backbone would be used primarily to connect the NSF supercomputer centers. The draft solicitation had indicated that the new network would continue to be a "general purpose" backbone, serving a large number of sites and carrying both routine and high-end traffic. By limiting the scope of the backbone, NSF's new approach would require more routine services to be obtained from commercial providers.

SOURCES: Robert Aiken et al., "NSF Implementation Plan for Interagency Interim NREN," May 1992; National Science Foundation, "Project Development Plan: Continuation and Enhancement of NSFNET Backbone Services," November 1991; National Science Foundation, "Network Access Point Manager/Routing Authority and Very High Speed Backbone Network Services Provider for NSFNET and the NREN Program: Program Solicitation," June 1992; Ellen Hoffman, "NSFNET Backbone Service Restructured," *Link Letter*, vol. 5, No. 3, November 1992, p. 1; Douglas E. van Houweling, Merit Network Inc., testimony at hearings before the House Subcommittee on Science, Mar. 12, 1992, pp. 36-41; Stephen S. Wolff, Assistant Director, Directorate for Computer and Information Science and Engineering, National Science Foundation, testimony at hearings before the House Subcommittee on Science, Mar. 12, 1992, pp. 133-138, pp. 148-156; William L. Schrader, President and CEO, Performance Systems International, Inc., testimony at hearings before the House Subcommittee on Science, Mar. 12, 1992, pp. 87-98; Ellen Messmer, "NSF Changes Course On Its Internet Plan," *Network World*, vol. 9, No. 51, Dec. 21, 1992, p. 1; Office of Inspector General, National Science Foundation, "Review of NSFNET," Mar. 23, 1993.

studied by a large number of research programs in addition to the testbeds. The issues affecting the deployment of these technologies in commercial settings are mainly concerned with trading the costs associated with the existing infrastructure against the potential of future markets for the new technologies. Regulatory and economic factors affecting the pace of deployment are beyond the scope of this background paper.

The involvement of the carriers in the testbeds was an important result of the visibility afforded by the HPCC program and the Corporation for National Research Initiatives' organization. All three major interexchange carriers and most of the Regional Bell Operating Companies are involved. The focus on ATM-related issues serves to provide experience with the construction of these networks and demonstrate their feasibility on a significant scale. Despite the carriers' stated commitment to ATM, the degree to which the transition to ATM represents a true paradigm shift for the telecommunications industry should not be underestimated. The testbeds will have served to help advance the carriers beyond the stage of standards-setting, component development, and small-scale experiments. There are many who believe that a nationwide gigabit network is not possible without basing it on the ongoing investments of the carriers.[14]

The testbeds may also be helping to provide input to the ATM standards process. Currently, there is some concern in the telecommunications industry that elements of ATM are being standardized before there is sufficient understanding of the tradeoffs. In particular, there is uncertainty about the best way to control the traffic in ATM networks, a key component in the use of ATM to support integrated services. The testbeds will provide experience with real traffic, due to the involvement of applications researchers. The academic researchers are also contributing to the solution of these problems; while algorithms for the control of packet networks are a longstanding topic of theoretical research, the testbeds may serve to focus the work of academic researchers on topics of concern to industry to a greater extent.

[14] Vinton G. Cerf, "Another Reading of the NREN Legislation," *Telecommunications*, vol. 25, No. 11, November 1991, p. 29.

Index

A
Access link, Internet, 29
Advanced Research Projects Agency
  administration of testbeds, 8, 55
  Arpanet development, 6
  Internet, 15
Algorithm, fast packet network, 62
American National Standards Institute, 34
Analog network, digital network, contrasted, 2-3
ANSI. *See* American National Standards Institute
Application, defined, 21
Arpanet, Internet evolution from, 6
Asymmetric Digital Subscriber Line, 47
Asynchronous Transfer Mode (ATM) switching technology, 13, 43-46, 49
  future applications, 68-69, 74
ATM. *See* Asynchronous Transfer Mode (ATM)
AT&T, contribution to testbed research, 11
AURORA network, 54, 59

B
"Backbone network"
  Federal agency future direction, 67-70, 72
  Internet, 15-17
Bandwidth
  of agency networks, used presently, 12
  defined, 3
  link, multiplexing, 28
  requirement, broadband network technology, 39, 41
  shortage of, Internet, 22
BARRNET, 17
Bay Area Regional Research Network. *See* BARRNET
Bit, defined, 2
BLANCA network, 56
Broadband Integrated Services Digital Network (B-ISDN), 46, 56-57
Broadband network technology, 35-49
  Asymmetric Digital Subscriber Line, 47
  Asynchronous Transfer Mode (ATM) switching technology, 43-46, 49
  bandwidth requirement, 39, 41
  classroom utilization, 36-37
  component development, 44-47
  compression technology, 41
  computer development, 46-47
  distributed computing, 37-39
  flexibility of, 41-42
  High-bit-rate Digital Subscriber Line, 47
  high-definition television signal, 41
  Integrated Services Digital Network, 47-48
  interactive visualization, 36, 38
  Internet application, 47-49
  multimedia applications, 36-37
  National Research and Education Network research, objective of, 6
  network requirements, and applications, 39-42
  optical fiber, 45-46
  overview of, 2-5
  public network, 49
  simulation, 36-39
  switch development, 46
  video, 36-38
Bulletin board, Internet, 21

C
Cable television network, testbed research technology application to, 71, 74
Campus network, and Internet, 17
CASA network, 58
Circuit switch, 6, 29, 32
Classroom, broadband network technology utilization, 36-38
Climate change, and network research, 1
Coaxial cable, computer link, 5
Color representation, bit patterning, 2
Commercial availability, of broadband network transmission equipment, 10
Commercial provider, Internet service, 19
Commercialism, of National Research and Education Network, 7
Compression technology, broadband network technology, 41
Computer network. *See* Network
Contractor, for National Research and Education Network, 7-8
Copper link, computer network, 5
Corporate network, Internet, 17
Corporation for National Research Initiatives, 8
Cost, network research component, of High Performance Computing and Communications program, 1-2

D
DARTnet, Internet, 15-16
Data, pattern of electronic signal, 2
Delay, in testbed operation, 10
Department of Energy. *See* U.S. Department of Energy
Digital computer, signal, 2
Digital network
  analog network, contrasted, 2-3
  defined, 2
Distributed computing
  broadband network technology, 37-39
  Internet, 22
Distributed supercomputer, 38, 40, 62-63
DOE. *See* U.S. Department of Energy

E
E-mail, Internet, 21
Electronic mail. *See* E-mail
Electronic signal, digital computer, 2
ESnet, Internet, 15-16
Exchange Carriers Standards Association, 34

F
Fast packet switching, 42-45
  algorithm, 62
  Asynchronous Transfer Mode (ATM) switching technology, 43-46
  for broadband network, 10
  Packet Transfer Mode (PTM), 43-44
Federal Internet Exchange (FIX), Internet, 17
Fiber optic link, 5
  transmission standards, Synchronous Optical Network, 28-29
File Transfer Protocol, Internet, 21
Frame Relay, 34
FTP. *See* File Transfer Protocol
Funding, testbed
  under Corporation for National Research Initiatives agreement, 8
  overview, 11

G
General Accounting Office, protest to, by National Research and Education Network contract bidder, 7-8
Gigabit research
  AURORA, 54, 59
  BLANCA, 56-57
  Broadband Integrated Services Digital Network (B-ISDN), 56-57
  CASA, 58
  MAGIC, 61
  network
    fiber optic link, 5
    overview, 4
    technology, progress of, 8, 10
  objectives of, 52-56
  standard, High Performance Parallel Interface, 53
  supercomputing research, 62-63
  Synchronous Optical Network (SONET), 65, 67
  testbed design, 53-55
  testbed progress in, 56-63
  VISTAnet, 60, 62
Government funding, of testbeds, 11
"Grand Challenges", network research and, 1

H
High-bit-rate Digital Subscriber Line, broadband network technology, 47
High-definition television signal, compression technology, 41
High-Performance Computing Act of 1991, legislative purpose, 1
High Performance Computing and Communications program, network research component
  gigabit network technology research, 51-63
  overview of, 1-13
High Performance Computing and High Speed Networking Applications Act of 1993, 21
High Performance Parallel Interface, gigabit-per-second standard, 53
HIPPI. *See* High Performance Parallel Interface

I
Image, bit patterning, 2
Industry investment, in testbeds, 11
Information handling, by computer network, overview, 2
Information pattern, voltage, 2
Integrated Services Digital Network broadband network technology application, 47-48
Interactive visualization, broadband network technology, 36, 38
Interagency Interim National Research and Education Network, 20
Interface, with Internet, 22
International Telecommunications Union, standards, establishment of, 34
International Telegraph and Telephone Consultative Committee, standards, establishment of, 34
Internet, 15-34
  access link, 29
  Advanced Research Projects Agency, 15
  applications, 21-24
  "backbone" networks, 15-17
  bandwidth, shortage of, 22
  BARRNET, 17
  broadband network technology application, 47-49
  bulletin board, 21
  campus network and, 17
  commercial providers, 19
  composition of, overview, 6-7
  corporate network and, 17
  DARTnet, 15-16
  distributed computing, 22
  e-mail, 21
  ESnet, 15-16
  Federal Internet Exchange (FIX), 17
  File Transfer Protocol, 21
  Frame Relay, 34
  Interagency Interim National Research and Education Network, 20
  LAN, 22, 24
  "massively parallel" computer, 24-25
  modem, 29
  National Science Foundation, 18
  network components, 26-31
    computer, 27-28
    link, 28-29
    switch, 29-32
  network technology, and applications, 22-24
  newsgroup, 21
  NSFNET, 15-16
    backbone, 19
    hierarchical structure, 18
  NSI, 15-16
  NYSERNet, 17
  packet switching, 24-26, 32
  private sector opportunity, 18-19
  protocols, 24-26
  public switched network, 31-34
  regional network and, 17-18
  sound, limitations of, 22
  Switched Multimegabit Data Service, 34
  technology used in, 21-34
  telephone company, 28-29, 31-34
  Telnet, 21
  testbed research application, 71
  transmission rates, 28-29
  transparent application version, 22
  TWBnet, 15-16
  users of, increase in number of, 12-13
  video, limitations of, 22-23
  WAN, 22, 24
Internet Activities Board, standards, 34, 70
Internet Engineering Task Force, standards, 34, 70
Internetworking, National Research and Education Network, 69-70, 72-73

L
LAN, Internet, 22, 24
Link, Internet, 28-29
Local area network. *See* LAN

M
MAGIC network, 61
"Massively parallel" computer, 24-25
MCI, contribution to testbed research, 11
Modem, Internet, 29
Multimedia applications, of broadband networks, 3, 36-37
Multiplexing, bandwidth link, 28

N
Narrowband network, defined, 3
National Aeronautics and Space Administration
  NASA Science Internet. *See* NSI
National Research and Education Network
  commercialism of, 7
  internetworking, 69-70, 72-73
  objective of, 5-8
  scope of, 7
  testbed research application, 65-70
National Science Foundation, Internet, 18
National Science Foundation Network. *See* NSFNET
NECTAR network, 59
Network, components of, 3, 5
  Internet, 26-31
Network design, broadband network, 3
Network technology, Internet applications, 22-24
New York State Education and Research Network. *See* NYSERNet
Newsgroup, Internet, 21
NREN. *See* National Research and Education Network
NSFNET
  future developments, 65-67, 70
  Internet, 15-16
    backbone, 19
    hierarchical structure, 18
  testbed research application, 65-67
NSI
  future development, 65-66, 68
  Internet, 15-16
NYSERNet, 17

O
Optical fiber, 45-46
  link, of broadband network, 10
  signal, digital computer, 2

P
Packet switching, 6
  Internet, 29, 31-32
Packet Transfer Mode (PTM), fast packet switching, 43-44
Private network
  establishment of, 33-34
  Federal agency, future development, 66
Private sector opportunity, Internet, 18-19
Procurement procedure, and National Research and Education Network, 7-8
Protocols, Internet, 24-26
Prototype gigabit network. *See* testbed
Public network
  broadband network technology, 49
  Internet, 31-34

R
Regional Bell Operating Company, contribution to testbed research, 11
Regional network, and Internet, 17-18
Remote login. *See* Telnet
Router. *See* Packet switch

S
Shade, bit patterning, 2
Simulation
  broadband network technology, 36-39
  and visualization, 6
SONET. *See* Synchronous Optical Network
Sound, Internet limitations, 22
Sprint, contribution to testbed research, 11
Standards, establishment of, for Internet, 34
Supercomputer
  distributed, 38, 40, 62-63
  and National Research and Education Network, 8
  and visualization, 6
Switch
  circuit, packet, contrasted, 29, 32
  development, broadband network technology, 46
  future application of testbed research, 68

Internet, 29-32. *See also* Packet switching
technology. *See also* Circuit switch, Fast packet switching, Packet switching
Asynchronous Transfer Mode (ATM), 13
overview, 35
Switched Multimegabit Data Service, 34
Synchronous Optical Network, 58, 65, 67
transmission standards, fiber optic link, 28-29

T
Telecommunications carriers, testbed research, contribution to, 11
Telephone company
Internet, 28-29, 31-34
link lease from, 28
Telnet, Internet, 21
Testbed
concept of, 10-11
design, gigabit research, 53-55
gigabit, National Research and Education Network, 8-9
organization, gigabit research, 55-56
Testbed research
advantages of, 11
application, National Research and Education Network, 65-70
application of, 65-74
Asynchronous Transfer Mode (ATM), 13
National Research and Education Network application, 12-13
Transmission
rates, Internet, 28-29
standards, fiber optic link, Synchronous Optical Network, 28-29
technology, future application of testbed research, 67-68
Transparent application version, Internet, 22
TWBnet, Internet, 15-16
Twisted pair cable, computer link, 5

U
U.S. Department of Energy, testbed research application, 65-66
User community, Internet, 6-7

V
Video
bit patterning, 2
broadband network technology, 36-37
Internet limitations, 22-23
Videoconferencing system, bit transmission, 3
VISTAnet, 60, 62
Visualization, by gigabit National Research and Education Network, 6
Voltage, information pattern, 2

W
WAN, Internet, 22, 24
Wide area network. *See* WAN

PLEASE PHOTOCOPY

Superintendent of Documents **Publications** Order Form

Order Processing Code:
*7110

P3
Telephone orders (202) 783-3238
To fax your orders (202) 512-2250
Charge your order.
It's Easy!

☐ **YES**, please send me the following:

_____ copies of ***Advanced Network Technology—Background Paper (88 pages)***,
S/N 052-003-01326-6 at $5.00 each.

The total cost of my order is $_____. International customers please add 25%. Prices include regular domestic postage and handling and are subject to change.

(Company or Personal Name) (Please type or print)

(Additional address/attention line)

(Street address)

(City, State, ZIP Code)

(Daytime phone including area code)

(Purchase Order No.)

Please Choose Method of Payment:

☐ Check Payable to the Superintendent of Documents
☐ GPO Deposit Account ☐☐☐☐☐☐☐ – ☐
☐ VISA or MasterCard Account

☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐

☐☐☐☐ (Credit card expiration date)

Thank you for your order!

(Authorizing Signature) (6/93)

May we make your name/address available to other mailers? YES ☐ NO ☐

Mail To: New Orders, Superintendent of Documents, P.O. Box 371954, Pittsburgh, PA 15250-7954

Office of Technology Assessment

The Office of Technology Assessment (OTA) was created by the Congress in 1972 as an analytical arm of Congress. OTA's basic function is to help Congressional policymakers anticipate and plan for the consequences of technological changes and to examine the many ways—expected and unexpected—in which technologies affect people's lives. The assessment of technologies calls for exploration of the physical, biological, economic, social, and political impacts that can result from applications of scientific knowledge. OTA provides Congress with independent and timely information about the potential effects—both beneficial and harmful—of technological applications.

Requests for studies are made by chairmen and ranking minority members of standing committees of the House of Representatives or Senate, by the Technology Assessment Board, the governing body of OTA, or by the Director of OTA in consultation with the Board.

The Technology Assessment Board is composed of six members of the House, six members of the Senate, and the OTA Director, who is a nonvoting member.

OTA has studied many ways to improve productivity, energy and materials conservation, technology, and employment; national security and international affairs; food and renewable resources; health; telecommunications and computer technology; [illegible]